

# No More Throwing Darts at the Wall: Developing Fair Handicaps for Darts using a Markov Decision Process

Other Sports  
Paper ID: 193939

## 1. Introduction

Handicap systems are commonly used in sports to create a competitive balance between players with different skill levels. They operate by providing some advantage to the weaker competitor to improve their chances of winning against the stronger player. Handicaps are widely used in many sports, most popularly in golf. Beyond creating competitive balance, handicaps can provide insights into other applications including player rankings and sports betting.

We focus on the popular sport of darts, where players throw darts at a circular board with numbered regions to score points. Each player starts with a score of 501 with the goal of reaching zero in as few throws as possible. Within darts, games are typically handicapped by providing the weaker player with an initial head start [1]. However, the head-start values are calculated using heuristics that are not proven to truly balance competition.

We seek to improve upon these heuristic methods as well as develop a novel alternative handicap system. For example, consider the possibility of a “credit” dart that the weaker player may use to “hit” a desired region of the board with a probability of one. Players may use these credits dynamically throughout the game, allowing them to address imbalances which may arise at any stage. An appropriate handicap system could afford the weaker player the exact number of credits that would create a competitive balance.

In this paper, we develop a framework to model the game of darts as a Markov Decision Process (MDP) with a dynamic credit handicap system. This framework is capable of modeling existing head-start handicaps as well as dynamic credits. We define the states of the MDP as the current score and number of available credits and solve for the optimal aiming location (i.e., target) for each state. A player can choose from hundreds of minute targets on the board. Players also possess inherent *execution error* which models the probability of where the dart will actually land given an intended target. We model execution error using bivariate Gaussian distributions derived from professional player data. We use this framework to study players of different skill levels by systematically perturbing the covariance matrices  $\Sigma$  of the Gaussian distributions, where larger  $\Sigma$  values represent less skilled players.

In summary, this paper makes the following contributions.

1. We develop and solve an MDP model for the game of darts with a dynamic credit-based handicap system, tuned for players with differing skill levels.
2. We rigorously show that the current heuristic head-start handicap does *not* balance competition. In response, we develop two optimization-based handicaps: an optimized head-start and a novel dynamic credit handicap.

3. Finally, we use simulation to demonstrate that dynamic credit handicaps create true fairness, which is very difficult (and sometimes impossible) to achieve with a head-start.

We believe that the creation of these fair handicap systems can benefit the growing darts community. Indeed, the popularity of darts in Europe has recently seen a substantial increase [2]. Moreover, it is estimated that around 17 million people play darts in the United States [3]. Standardized, fair handicaps can be used to balance competition and increase enjoyment for players of all levels from professionals and amateurs playing in tournaments to novices playing for fun at the pub.

## 2. Related Literature

This paper builds on three areas of prior research: (1) modeling the optimal strategy in darts (2) modeling skill level and execution error in sports and (3) developing rigorous handicapping frameworks in sports.

Attempts to model the optimal strategy for darts began in the 1980s when Kohler (1982) used dynamic programming to model the game, albeit with several simplifying assumptions [4]. More recently, Haugh and Wang (2022) developed a sophisticated approach to solve for a player's optimal darts strategy while also considering their opponent; they used a zero-sum game approach in combination with an underlying MDP formulation [2]. They demonstrated that, for professionals, considering one's opponent led to a marginal increase in win probability of just 0.2-0.6%.

There has also been a growing interest in modeling execution error and skill level, both in darts and in sports more broadly. Within darts, Tibshirani et al. (2011) developed a methodology to model a player's execution error as a bivariate Gaussian distribution based on throw data using the Expectation Maximization (EM) algorithm [5]. Haugh and Wang (2022) built on this work by using a novel data set of dart-throws by 16 professional darts players to fit player-specific distributions for various targets on the board [2]. Finally, Chan et al. (2022) explored how the optimal strategy in tennis relates to a player's "level" of execution error [6]. They denote this level as  $\epsilon$ , a positive multiplicative factor that scales the covariance matrix of the underlying execution error distributions, with higher  $\epsilon$  values corresponding to a higher execution error.


Finally, to date, limited formal literature has explored rigorous handicap systems for sports. Bingham and Swartz (2000) evaluated traditional handicaps for golf, highlighting the existence of unfairness and proposing a statistically fairer alternative [7]. More recently, Chan and Singal (2016) developed an MDP-based handicap system for tennis using a dynamic credit system [8]. However, no literature currently discusses handicaps in darts.

This paper bridges prior work from the above three areas to develop a handicapping framework for the game of darts. Specifically, we develop a novel MDP formulation that can incorporate various levels of skill (i.e., execution error) and subsequently evaluate various handicap systems.

## 3. Background on the Game of Darts

### 3.1 The Rules of Darts

This paper will focus on 501 darts which is the most common darts variant. The game is played between two players, each with a starting score of 501. Players take turns throwing three darts at a time. Each throw is worth a certain number of points depending on where the dart lands on the board. The cumulative points gained across all three throws at the end of the turn are then subtracted from the player's score at the beginning of the turn to give their new score.



**Figure 1:** Layout of a Standard Dartboard

For each throw, points are allocated based on where the dart lands on the board (visualized in Figure 1). First, hitting the outer ring is called a “double” (“D”) and will provide double the points specified for that slice of the board (i.e., the number beside a given slice). Second, hitting the inner ring is called a “triple” (“T”) and will provide triple the score specified for that slice of the board. Third, hitting a non-ring area of the board is called a “single” (“S”) and will provide the exact score specified for that slice of the board. For example, Figure 1 highlights hitting the S20, D20 and T20 regions on the board, which result in receiving 20, 40, and 60 points, respectively. Moreover, hitting the outside of the center circle is called a “single-bullseye” (“SB”) and will provide 25 points. Finally, hitting the inside of the center circle is called a “double-bullseye” (“DB”) and will provide 50 points.

To win (or “check out”), a player must hit a double that brings their score to exactly zero. Otherwise – if a player’s throw brings them to a negative score, a score of one, or reaches zero but is not a double – they “go bust”. This means that their turn is invalidated and their score reverts to its value at the start of the turn.

### 3.2 Existing Handicaps in Darts

Several unofficial handicap systems are used in darts, with the most common involving a head-start for the weaker player. The determination of this head-start value relies on the main statistic measuring player skill in 501 darts: points per dart (PPD) [1]. PPD is calculated as follows:

$$PPD = \frac{\text{Total Points Scored}}{\text{Total Number of Darts Thrown}} \quad (1)$$

This can be based on an individual game, a series of games, or a player's lifetime. To determine the head start, it is common to use a simple heuristic: the ratio between the two players' PPD values.

Our framework allows us to understand and then improve this current handicap in three ways. First, we mathematically demonstrate the imbalance that persists with the current head-start handicap based on PPD. Second, our framework provides a more precise representation of player skill, distinguishing between execution error and policy. In contrast, PPD-based handicaps may not differentiate between the two; a player may have a lower PPD because they take a more conservative policy regardless of their throw accuracy. Lastly, we introduce a dynamic credit handicap which offers flexibility for implementation at any game stage, unlike existing handicaps assigned solely at the start. This flexibility allows us to identify crucial (or "clutch") throws for players of varying skill levels.

## 4. Markov Decision Process Model

In this section, we describe the various components of our MDP formulation for the game of darts with dynamic credits.

### 4.1 States


We define each state as  $s = (\theta, \beta, i, u)$ , where  $\theta \in \{0, \dots, 501\}$  is the current score,  $\beta \in \{0, \dots, 9\}$  is the number of remaining credits,  $i \in \{1, 2, 3\}$  is the number of throws remaining in the current turn and  $u \in \{0, \dots, 120\}$  is the number of points gained so far in the current turn. We only solve for up to nine credits because the shortest path from a starting score of 501 to 0 is nine throws, which will always be achieved when playing optimally with nine or more perfect-accuracy credits. An absorbing state is reached when a player wins the game, denoted  $s_{end} = (0, \beta, i, u)$  for all values of  $\beta, i, u \geq 0$ . This results in over 1.8 million distinct states. We highlight that setting  $\beta = 0$  allows us to easily model the original game of darts with no credits.

### 4.2 Actions

The action set  $A = \{A_{throw} \cup A_{credit}\}$  is the union of two subsets. First, the actions in the set  $A_{throw}$  correspond to 729 targets on the board that a player may choose to aim for if they are not using a credit (see Figure 2).

Second, the actions in set  $A_{credit}$  correspond to the 62 regions that a player may choose to "hit" with perfect execution if they use a credit. These 62 regions include twenty singles, twenty doubles, twenty triples, the single bullseye and the double bullseye.

We define a *policy* as a function  $\pi: S \rightarrow A$  which maps all states to an action that should be chosen in that state. By solving the MDP, we will determine the *optimal policy*, that is, the optimal strategy a player should use in darts.



**Figure 2:** Set of targets that a player may aim for if using a throw action.

### 4.3 Costs

There is a cost  $c(s') = 1$ , for every case where  $i' = 3$ , indicating that a turn has been completed. Consequently, the value function of our MDP will correspond to the expected number of turns needed to complete the game from a given state.

### 4.4 Transition Probabilities

The transition probabilities refer to the probability of starting in state  $s$ , choosing action  $a$  and transitioning to state  $s'$ , denoted as  $P(s'|s, a)$ . Within our formulation, the stochasticity comes from the distribution  $P(z|a)$ , where  $z$  is the resulting dart score (e.g., D1 or S12) given the chosen action  $a$ . We describe the transition dynamics in more details in the following subsections.

#### 4.4.1 Transition Distributions

The transition probabilities can be broken down into two distinct distributions depending on the type of action. First, if  $a \in A_{credit}$  (i.e., a player uses a handicap credit), their transition probability is deterministic: they will receive the desired number of points for their chosen move with a probability of one. Specifically,  $P(z|a) = 1$  where  $z$  corresponds to the dart score of the chosen credit region and  $P(z|a) = 0$  otherwise.

Second, if  $a \in A_{throw}$  (i.e., a player decides to throw), their transition probability is stochastic: the number of points they receive will depend on the dart's realized landing location (i.e., result) given a particular target. Consequently, we utilize fitted skill models for darts players, which has been the focus of many past studies [5, 9, 2]. Since this is not the main focus of our paper, we opt to use the fitted skill models from the most recent study, which fits bi-variate Gaussian distributions to throw data from professional darts players. These distributions are of the form  $\mathcal{N}(\mu, \Sigma)$ . Next, we divide the board into discrete 2mm-by-2mm pixels and integrate the probability mass of  $\mathcal{N}(\mu, \Sigma)$  over each result to compute the distribution  $P(z|a)$ .

#### 4.4.2 Transition Dynamics

After determining the probabilities  $P(z|a)$ , we must define the transition dynamics of 501 darts and the credit handicap system to determine the future state  $s' = (\theta, \beta, i, u)$  given the action result  $z$ . This is denoted by the state transition function which we call  $f$  (defined in Equation 2) where  $s' = f(s, z)$ . In this function, we introduce  $h(z)$  which represents the number of points gained from the dart landing region  $z$ . For example,  $z = T20$  will have  $h(z) = 60$ .

The state transition function consists of eight cases. The first four cases correspond to the stochastic transitions when a player uses a throw action. The first case corresponds to the player winning the game. This occurs when the final throw value brings the score to zero and was a double. The second case describes a player going bust, which results in the final throw becoming invalidated and their score returning to what it was at the start of their turn. The third case corresponds to a regular throw within their current turn. The added value is stored in the variable  $u$  and their remaining throws in the turn  $i$  increments down by one. The fourth case corresponds to the player making a regular throw to end their turn. The value stored in  $u$  across the turn finally increments their actual score  $\theta$ ,  $u$  resets to a value of zero and they are given three throws for their next turn.

The second four cases correspond to the deterministic transitions when a player uses a handicap credit, and mirror the cases described above. However, these transitions are only possible when the current state has at least one handicap credit available ( $\beta > 1$ ) and always transition to a state with one fewer handicap credits ( $\beta' = \beta - 1$ ).


$$f((\theta, \beta, i, u), z) := \begin{cases} (0, \beta, 0, 0), & a \in A_{throw}, \theta = u + h(z), z \text{ is double} \\ (\theta, \beta, 3, 0), & a \in A_{throw}, \theta - (u + h(z)) \leq 1 \\ (\theta, \beta, i - 1, u + h(z)), & a \in A_{throw}, \theta > (u + h(z) + 1), i > 1 \\ (\theta - u - h(z), \beta, 3, 0), & a \in A_{throw}, \theta > (u + h(z) + 1), i = 1 \\ (0, \beta - 1, 0, 0), & \beta > 0, a \in A_{credit}, \theta = u + h(z), z \text{ is double} \\ (\theta, \beta - 1, 3, 0), & \beta > 0, a \in A_{credit}, \theta - (u + h(z)) \leq 1 \\ (\theta, \beta - 1, i - 1, u + h(z)), & \beta > 0, a \in A_{credit}, \theta > (u + h(z) + 1), i > 1 \\ (\theta - u - h(z), \beta - 1, 3, 0), & \beta > 0, a \in A_{credit}, \theta > (u + h(z) + 1), i = 1 \end{cases} \quad (2)$$

#### 4.4.3 Transitions for Different Skill Levels

In order to model varying skill levels, we can systematically perturb  $P(z|a)$ . For this process, we were inspired by the tennis execution error framework introduced by Chan et al. (2022) [6]. We use their execution error level parameter  $\epsilon$ , which acts as a scalar error multiple on the covariance matrices of our chosen average bi-variate Gaussian distributions, that is,  $\mathcal{N}(\mu, \epsilon\Sigma)$ . For ease of exposition, all of our analysis was based off of the underlying distributions for the professional player Michael Smith, who was chosen to represent the average professional.




The bi-variate Gaussian distribution when aiming at the double-bullseye for this “average player” with varying execution error levels  $\epsilon$  can be seen in Figure 3. Based on our formulation,  $\epsilon = 1$  corresponds to professional-level accuracy with execution error increasing with larger  $\epsilon$  values. As can be seen, players with near-professional accuracy are able to pinpoint the double bullseye whereas players with higher execution error will regularly miss and hit the single region.



**Figure 3:** Bi-variate Gaussian distributions for players with different execution error levels  $\epsilon$  when aiming at the double bullseye.

It should be noted that while our chosen player has a radially symmetric distribution around the double-bullseye, this will not be the case for all players. Our modelling framework can be easily extended to analyze players having distributions of different shapes. For example, the distribution shapes for three different professional players (Cullen, Chisnall and Price) are shown in Figure 4. Not only can the distributions be asymmetric, but they can have different angular orientations with respect to the axes.



**Figure 4:** Bi-variate Gaussian distribution shapes for three different professional players with  $\epsilon = 10$  to allow for easier visualization.

Our analysis will hold shape constant and instead focus on how the optimal policy changes with  $\epsilon$  by re-integrating  $P(z|a)$  for  $a \in A_{throw}$  and re-solving for the optimal policy given  $P_\epsilon(s' | s, a)$ .

## 4.5 Bellman Equation

We can combine the states, actions, costs and transition probabilities outlined above to construct an appropriate MDP Bellman equation to compute the value of using a given policy  $\pi$ . By construction, we can interpret  $V_\epsilon^\pi(s)$  as the expected number of turns required for a player with execution error  $\epsilon$  to check out from state  $s$  when using policy  $\pi$ .

$$V_\epsilon^\pi(s) = \sum_{z \in Z} c(f(s, z)) + P_\epsilon(z|\pi(s)) V_\epsilon^\pi(f(s, z)) \quad (3)$$

We can use the following minimization operators to compute the optimal value and optimal policy that minimizes the expected number of turns to check out from any state. We solve this model using a standard policy iteration algorithm.

$$V_\epsilon^*(s) = \min_{a \in A} \left\{ \sum_{z \in Z} c(f(s, z)) + P_\epsilon(z|a) V_\epsilon^*(f(s, z)) \right\}, \forall s \in S \quad (4)$$

$$\pi_\epsilon^*(s) = \operatorname{argmin}_{a \in A} \left\{ \sum_{z \in Z} c(f(s, z)) + P_\epsilon(z|a) V_\epsilon^*(f(s, z)) \right\}, \forall s \in S \quad (5)$$

Given the large sample state space, we are able to exploit the monotonic structure of darts to solve this more quickly as a dynamic program in practice.


## 5. Properties of Value Function and Optimal Policy

In this section, we analyze the optimal value and optimal policy of the MDP with varying levels of execution error and number of handicap credits.

### 5.1 Optimal Policy and Execution Error


Figure 5 illustrates how the distribution of optimal actions (i.e., targets) changes as a function of  $\epsilon$ , when no handicap credits are available. When  $\epsilon$  is low (e.g.,  $\epsilon = 1$ ), the optimal policy is concentrated at the T20, indicating that players should repeatedly target this region to reliably receive 60 points per throw (the most a single dart can achieve). This is continued until the end of the game when they are trying to check out. At this stage, if they have an even score, they will target the required double. If they have an odd score then they will target an odd single that will move them to an even score where they can finish with a double; for near-professional players, even-valued singles are rarely targeted. As  $\epsilon$  increases (e.g.,  $\epsilon = 9$ ), this policy shifts towards the T19 instead of the T20. This is likely because when a player misses a T19 (neighbored by 7 and 3), they receive more points than if they miss a T20 (neighbored by 5 and 1).





**Figure 5:** Distribution of optimal policy across all score states for different levels of execution error  $\epsilon$ .


As epsilon increases further (see Figure 6), the policy begins to constrict and targets a more compact region that is located towards the centre of the board. In an extreme case with  $\epsilon = 100$ , the optimal policy suggests to primarily target the centre of the board to avoid missing the board completely.



**Figure 6:** Bi-variate Gaussian distributions for double-bullseye target (top) and corresponding distribution of optimal policy (bottom) across all score states for larger levels of execution error  $\epsilon$ .

## 5.2 Optimal Value and Execution Error

Figure 7 shows the optimal value (i.e., expected number of turns to finish the game) of each current score, at the start of a turn, with different levels of  $\epsilon$  and no handicap credits. As expected, larger  $\epsilon$  values increase the expected number of turns needed. However, the rate of increase declines as  $\epsilon$  grows. For example, the difference in value between  $\epsilon = 1$  and  $\epsilon = 2$  is larger than the difference between  $\epsilon = 8$  and  $\epsilon = 9$ .




**Figure 7:** Expected number of turns to finish the game  $V_\epsilon^*$  as the execution error  $\epsilon$  changes.

For scores greater than 40, the expected number of turns to finish the game decreases linearly as the score decreases. However, as we approach the end of the game, this number begins to oscillate. As was mentioned in Section 5.1, this is driven by the need to end on a double. That is, any odd-numbered score below 40 must first target an odd-numbered region before it can subsequently target a double to check out. This is contrasted with even-numbered scores that can target a double immediately. We also note the sharp increase in expected number of turns near the end of the game. This is driven by the difficulty of hitting a specific double; there is a high probability of going bust or missing the board in this range. Both of these outcomes keep a player's score the same, leading to a cycle where a player is stuck in the same state until they finally check out. Moreover, as their score approaches zero, there are fewer segments of the board that a player can hit without going bust; in the extreme case, with a score of 2, any throw result other than a D1 will keep their score the same. In contrast, with a score of 40, even if a player misses the desired D20, there is a large area of the board that they can hit without going bust.

## 5.3 Optimal Policy and Value with Handicap Credits

Figure 9 shows the optimal value at the start of the game (i.e.,  $\theta = 501, i = 3, u = 0$ ) with different levels of  $\epsilon$  and number of dynamic credits  $\beta$ . By comparing the expected number of turns needed at the start of the game, we arrive at a single catch-all metric to compare different handicap frameworks. As expected, as available credits increase, the expected turns decrease. Note, the

expected number of turns approach a limit of 3, which corresponds to a perfect “9-dart finish” within those three turns; it is not possible to get from an initial score of 501 to a score of zero in less than 9 throws. We also observe that the first credit brings the most value, indicated by the steeper slope between  $\beta = 0$  and  $\beta = 1$ . This is because the optimal policy instructs players to always save their credits for the end of the game, which removes the cycle of missing the board or going bust. As  $\beta$  increases past 1, there is a steady linear decline in the number of expected turns to finish.



**Figure 9:** Expected number of turns to finish the game  $V_{\epsilon}^*(\theta = 501, \beta, i, u)$  as the execution error  $\epsilon$  and credits  $\beta$  change.


Going deeper, we are able to observe "break-even" points where two players with different  $\epsilon$  values can be given credits to equalize their number of expected turns. For example, a player with  $\epsilon = 1$  requires just under 6 turns to finish. This is equivalent to an  $\epsilon = 2$  player receiving approximately 3 credits, which can be easily obtained by drawing a horizontal line on Figure 9. This type of analysis can be used to understand how dynamic credits can create competitive balance between players with mismatched skill (see section 6.3).

## 5.4 Optimal Value and Heuristic Head-Start Handicaps

We now use this framework to examine the effectiveness of the heuristic headstart handicap system at equalizing competitive balance between mismatched players. We first make the assumption that all players will be using an optimal policy. We then divide 501 by the starting state  $V_{\epsilon}^*(\theta = 501, \beta = 0, i = 3, u = 0)$  to compute their optimal expected PPD (i.e., number of points divided by number of throws). Using PPD estimates for players with different  $\epsilon$  values, we are then able to calculate the weaker player's head start using the heuristic employed in practice [10].

$$\text{Weaker Player Starting Score} = 501 \left( \frac{PPD_{\text{weaker}}}{PPD_{\text{stronger}}} \right) \quad (6)$$

Figure 10 uses dots to indicate the weaker player's head start when competing against a player with  $\epsilon = 1$ . It is clear to see that these head starts still result in the weaker player requiring more turns to finish the game compared to the stronger player. For example, even with this heuristically computed head start, a player with  $\epsilon = 9$  would still require roughly two more turns to check out when faced against an  $\epsilon = 1$  player. We seek to address this imbalance with optimization based handicaps in the next section.




**Figure 10:** The expected number of turns to finish the game as execution error changes. The markers represent the state at which a player would start the game using the heuristic head-start handicap against a professional with  $\epsilon = 1$ .

## 6. Optimization-Based Handicaps

In this section, we use our MDP to propose two optimization-based handicaps: (i) an *optimized* head start handicap and (ii) a credit-based handicap. Both handicap systems seek to equalize the expected turns to check out from the start of the game. We note that these methods require estimating one's own value of  $\epsilon$ , which we describe in Appendix A.

### 6.1 Optimized Head-Start Handicap

As previously noted, the heuristic handicap still leaves the weaker player at a disadvantage. We can use our MDP to determine the appropriate level of head start that balances competition. Specifically, we take the  $V^*(\theta = 501, \beta = 0, i = 3, u = 0)$  of the stronger player and determine the starting score where the weaker player has the same number of expected turns to finish. Figure 11 (a) provides an illustration of this methodology and Figure 11(b) provides a lookup table handicap value.



(a) Markers represent the optimized head-start handicap needed against player with  $\epsilon = 1$ .


(b) Lookup table for the optimal head start needed for any two players.

**Figure 11:** Optimized head-start handicap methodology (a) and lookup table (b).

Head-start handicaps have the advantage of being highly intuitive and aligned with current heuristic methods. However, there is a notable shortcoming: as the difference in  $\epsilon$  gets very large, there could be an irreconcilable discrepancy caused by the difficulty at the end of the game. That is, if it takes a player with high  $\epsilon$  more throws to hit a double than it does for a player with low  $\epsilon$  to complete the entire game from the initial score of 501, then no head start will ever be able to achieve competitive balance. This limitation can be addressed by using dynamic credits.

## 6.2 Novel Dynamic Credit Handicap

This handicap system involves giving  $\beta$  credits to the weaker player at the start of the game to equalize the expected number of turns needed to finish. Specifically, we take the  $V^*(\theta = 501, \beta = 0, i = 3, u = 0)$  of the stronger player and determine the necessary number of credits where the weaker player has the same number of expected turns to finish. Figure 12(a) provides an illustration of this methodology and Figure 12(b) provides the lookup table containing number of credits the weaker player should be allotted. Fractional values can be interpreted as how many credits the weaker player should be afforded on average over multiple games. Indeed, in practice, most darts competitions include multiple games, making fractional credits possible to equalize expected number of turns.



(a) Number of credits needed against a player with  $\epsilon = 1$ .

(b) Lookup table for the number of credits needed for any two players.


**Figure 12:** Dynamic credit handicap methodology (a) and lookup table (b).

## 7. Evaluating Handicap Fairness in Practice

Thus far, we have developed handicaps that equalize the expected turns required for each player to check out. We now evaluate their ability to equalize *win probability* between players (i.e., true fairness). This was accomplished by simulating 10,000 games between each set of players, alternating the starter (to control for the starters advantage). We assumed that both players use their respective optimal policies and sample from their throw outcome distribution. The first player to check out was recorded as the winner. We then estimated the true win probabilities as the proportion of games that each player won.

Figure 13 shows the win probabilities for the two optimization-based handicap methods. We find that the optimized head-start handicap does not equalize win probabilities and actually creates a disadvantage for the stronger player. Meanwhile, the novel dynamic credit handicap is much more successful at equalizing win probabilities (i.e., values closer to 50%). We also confirmed that the heuristic head start methods currently used in practice do not create fair outcomes, affording an advantage to the stronger player (see Appendix B).



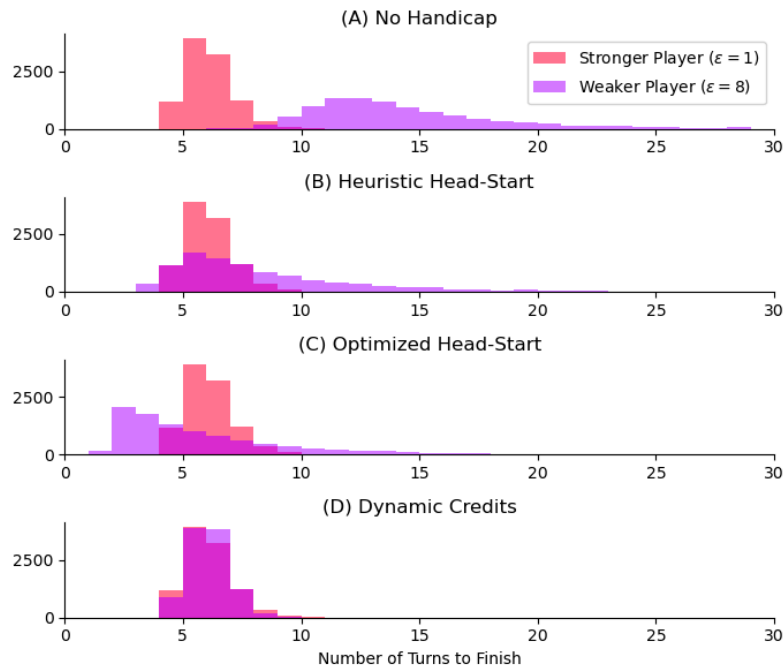


(a) Optimized head-start handicap.

(b) Novel dynamic credit handicap.

**Figure 13:** Probability that stronger player wins after 10,000 simulated matches.

Figure 14 shows a histogram of the number of turns to check out for two mismatched players with ( $\epsilon = 1$  and  $\epsilon = 8$ ) using different handicap approaches.



**Figure 14:** Number of turns to check out for two mismatched players for 10,000 simulated games.

We see that the weaker player's turns-to-check-out distribution has a higher mean, variance, and skew. Moreover, head-start handicaps merely shift the mean of the turns-to-check out distribution

to be earlier in the game, with the shape remaining the same. Therefore, equalizing the expected turns to check out will also equalize win probability *only if* the variance and skew are also equal for both players. Remarkably, this is exactly what the dynamic credit handicap achieves. By removing the difficulty at the end of the game, the turns-to-check out distribution for the weaker player is reshaped to mimic that of the stronger player, producing a close approximation of true fairness.

## 8. Conclusion

In this paper, we used a Markov Decision Process to develop a framework to handicap darts. We explored how this handicap system can be applied to equalize win probabilities for players possessing different levels of execution error. Our analysis shows that the main existing handicap system is imbalanced and that the most challenging component of the game is hitting a double to check out at the end. We then developed two optimization-based alternatives that equalize the expected turns for both players to win the game. Finally, we use simulation to evaluate the true fairness of each handicap and illustrate that our novel dynamic credit handicap produces the best approximation to true fairness.

There are several opportunities to extend this work including evaluation of new handicaps, deeper exploration of the dynamic credit handicap, and practical applications. First, this framework can be modified to study a range of handicap designs (e.g., credits could allow “do-over” throws instead of providing players with a guaranteed result). Second, more work can be done to investigate the effectiveness of dynamic credits at equalizing win probabilities by incorporating the opponent’s score into the state space, bringing the number of states to over one billion. Solving this scale of model is impractical but perhaps possible with enough compute power and programming that exploits the monotonic structure of darts. Solving a few one-off cases to test whether our simplified modeling approach can approximate this complicated MDP would provide a valuable benchmark. Finally, from a practical lens, we hope to partner with a local darts league to trial dynamic credit handicaps and see how they perform. In addition, our framework could be built into dart machines or embedded into an app to compute handicaps both for tournaments and casual matches. Beyond these darts-specific ideas, dynamic credits could also be applied and studied within other sports.

## References

- [1] J.W. Houriet, J.C. Nydick, D.L. Hedin, D.L. Aymar, T.R Horne, and R.H Wiles. "System for handicapping substitute or unranked players in a dart game match", 2000. U.S. Patent 6 076 021, Jun. 2000.
- [2] M. B. Haugh and C. Wang. "Play like the pros? Solving the game of darts as a dynamic zero sum game". *INFORMS Journal on Computing*, 2022. doi: 10.1287/ijoc.2022.1197. URL <https://arxiv.org/abs/2011.11031>.
- [3] Whitney McIntosh. "17 million americans play darts: so where are the world-class stars?", 2019. URL <https://www.theguardian.com/sport/2019/jun/26/darts-united-states-stars-growth>.
- [4] David Kohler. "Optimal strategies for the game of darts". *The Operational Research Society*, 1982. doi: 10.1057/jors.1982.191. URL <https://doi.org/10.1057/jors.1982.191>.
- [5] Ryan J. Tibshirani, Andrew Price, and Jonathan Taylor. "A statistician plays darts". 2011. doi: 10.1287/ijoc.2022.1197. URL <https://doi.org/10.1111/j.1467-985X.2010.00651.x>.
- [6] Timothy C. Y. Chan, Douglas S. Fearing, Craig Fernandes, and Stephanie Kovalchik. "A Markov process approach to untangling intention versus execution in tennis". *Journal of Quantitative Analysis in Sports*, 18(2):127–145, 2022. doi: doi:10.1515/jqas-2021-0077. URL <https://doi.org/10.1515/jqas-2021-0077>.
- [7] Derek Bingham and Tim Swartz. "Equitable handicapping in golf". *The American Statistician*, 54:170–177, 08 2000. doi: 10.1080/00031305.2000.10474541. URL <https://doi.org/10.1080/00031305.2000.10474541>.
- [8] Timothy C. Y. Chan and Raghav Singal. "A Markov decision process-based handicap system for tennis". *Journal of Quantitative Analysis in Sports*, 12(4):179–188, 2016. doi: doi:10.1515/jqas-2016-0057. URL <https://doi.org/10.1515/jqas-2016-0057>.
- [9] Thomas Miller and Christopher Archibald. "Monte Carlo skill estimation for darts". 2021 *IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8, 2021. doi: 10.1109/SSCI50451.2021.9659951.
- [10] Arachnid. "Handicapping", 2010. URL <http://www.arachnid360.com/wp-content/uploads/2010/03/Handicapping.pdf>.

## Appendix A. Estimating Execution Error Level

We describe how to estimate the skill level of each player. We acknowledge that more sophisticated methods to determine a player's skill level exist, but we focus on developing a method that is easy to interpret and implement. That is, we first ask players to throw 30 darts aimed at the double-bullseye and then use Table 1 to find their assigned execution error level  $\epsilon$ .


**Table 1:** Lookup table to determine a player's  $\epsilon$  value based on a sample of 30 throws. Columns indicate number of expected throws ending up in each region.

$\epsilon$	DB	SB	Other
1	10	16	4
2	5	14	11
3	4	11	15
4	3	9	18
5	2	8	20
6	2	7	21
7	2	6	22
8	1	6	23
9	1	5	24

Table 1 is derived by integrating the Gaussian distributions for varying  $\epsilon$  when they aim for the centre of the board. We then discretize the proportion of mass that falls in the DB, SB, and elsewhere. Multiplying these proportions by 30 provides the expected number of darts that should land in each region. We can accordingly assign a player their  $\epsilon$  in an interpretable manner.

## Appendix B. Heuristic Head Start Simulation Results

Figure 15 shows the win probability of the stronger player with the currently-used heuristic handicap. We see that this methodology does not create fair match outcomes and provides an advantage to the stronger player. For example, when a player with epsilon 1 plays against a handicapped epsilon 9 player, the stronger player still has a 70% chance of winning.



**Figure 15:** Probability that stronger player wins after 10,000 simulated matches using the heuristic head start handicap employed in practice.