An Outline of a Basic TCG Sportsbook Model Using Microsoft Excel, R, and Bayesian Inference

Nicholas A. Beaver

NicholasABeaver.com

7 March 2025

ABSTRACT

This paper outlines a baseline sportsbook model for wagering on competitive trading card games (TCGs) such as *Magic: The Gathering*, *Pokémon*, and *Yu-Gi-Oh!*. Using the Glicko-2 rating system, a logistic distribution to assess deck strength's impact on competitor win probability, Monte Carlo simulations, and Bayesian inference, we evaluate the feasibility and profitability of such a sportsbook from the perspective of the operator. A simulated world championship tournament, repeated over 10,000 iterations, representing 310,000 simulated matches, assesses expected profitability as measured by gross gaming revenue (GGR). Findings indicate a 75.7% probability of achieving a GGR margin above 5%, suggesting that a TCG sportsbook could be a viable business model.

Beyond theoretical validation, this paper discusses key challenges, including data availability and regulatory considerations. Future research should focus on refining the model with empirical data, implementing real-time probability and pricing adjustments, conducting necessary market research, and addressing legal and operational hurdles. This study provides a foundational approach to TCG wagering within the broader "esports" betting landscape upon which future work might be built.

Keywords: Bayesian inference, sports betting, trading card games, Glicko-2, Microsoft Excel, R programming, gambling models, financial modelling, sportsbook profitability

I. INTRODUCTION

In recent years, "esports" betting has emerged as a new space within the sports betting market, with odds now available for competitive videogames such as *League of Legends* and *Call of Duty* [1]. Near to the videogame esports now being played fiercely across the globe are trading card games (TCG's) such as *Magic: The Gathering, Pokémon,* and *Yu-Gi-Oh!*. These games claim millions of loyal fans and active competitive environments with competitors from all over the world competing in a multitude of regional and international tournaments. During *Magic's* 2025 Pro Tour, $500,000 in prize money will be offered to top players [2]. *Pokémon* plans to offer $2 million in prize support for the 2025 competitive year [3]. Despite this, no sportsbook currently offers odds on these events. TCG's represent potentially untapped opportunities for sportsbook operators.

This paper proposes a game-agnostic sportsbook model for TCG betting to answer the question "What would a sportsbook operator have to do to build a model to offer bet prices to customers for TCG events and how profitable would such a model be?." Thus, we only evaluate the feasibility and profitability from the perspective of "the house," or sportsbook operator. The model integrates key competitive factors—competitor skill and deck strength—to estimate competitor win probabilities and to consequently price odds for bettors.

This study focuses on three key research questions:

1. Can a sportsbook consistently generate profit by offering odds on TCG events?
2. How can competitor skill and deck strength be modelled to determine win probabilities for competitors and thus, price bet odds?
3. What key challenges must be addressed before real-world implementation?

To answer these questions, we construct a model that takes a synthetic dataset representing a "tournament season" consisting of five-hundred unique competitors engaged in 12,000 competitive matches to calculate competitor skill ratings and deck strength using the Glicko-2 rating system. We then use these ratings to set up a "world tournament" during which the top thirty-two competitors face off in a five-round single-elimination event. We compute win probabilities and betting odds, as well as the money taken in from and paid out to bettors, to calculate the net financial gain or loss for the house during each match of the world tournament. We simulate this world tournament over 10,000 iterations and use this dataset to draw conclusions about the effectiveness of the model. We then use this dataset to analyze expected profitability as measured by gross gaming revenue (GGR) margin. The final dataset contains 310,000 individual matches between competitors.

While this study demonstrates the theoretical feasibility of a TCG sportsbook, it does not address real-world bettor behavior, regulatory approval, or operational challenges, all of which require further research. The goal of this paper is to establish a foundational model and assess its potential viability. Future work should refine this framework using real-world tournament data, market research, and the implementation of real-time data feeds and automated computation of win probabilities and odds prices based on available data.

All work for this model is done in Microsoft Excel and with the open source `R` programming language. All files and code used for the model are available on the project website [4]. This paper represents an initial exploration of the concept. Feedback, critique, and suggestions for further refinement are welcome.

II. THEORY

A. Skill as a Gaussian Distribution

A competitor in a trading card game match engages in a game of skill against another competitor. These matches are subject to some level of randomness (statistical "noise"). Each competitor brings with him or her a deck, constructed beforehand, that uses a unique set of cards to bring to the gaming table a strategy designed to defeat opponents and secure a victory within the confines of the game's rules.

To summarize our belief in how competitor skill should be accounted for in our model, we subscribe our competitive skill theory to the following mathematical expression:

$$Skill_i \sim \mathcal{N}(skill_i^t, \sigma) \tag{1}$$

That is, each competitor *i,* has "skill" at time *t,* taking the form of a Gaussian distribution. This is a slight modification from the assessment of real-valued skill from Minka et al as used in Microsoft's competitive videogames [5]. Competitor skill taking the form of a Gaussian distribution means that we infer the real value of a competitor's skill, call him or her $Competitor_i$, within a range of probable skill level, as shown in Figure 1. As $Competitor_i$ continues to play the game (that is, as the value of *t* increases), the deviation of his or her skill (that is, σ) should become smaller and smaller, representing our increased confidence in the skill rating.

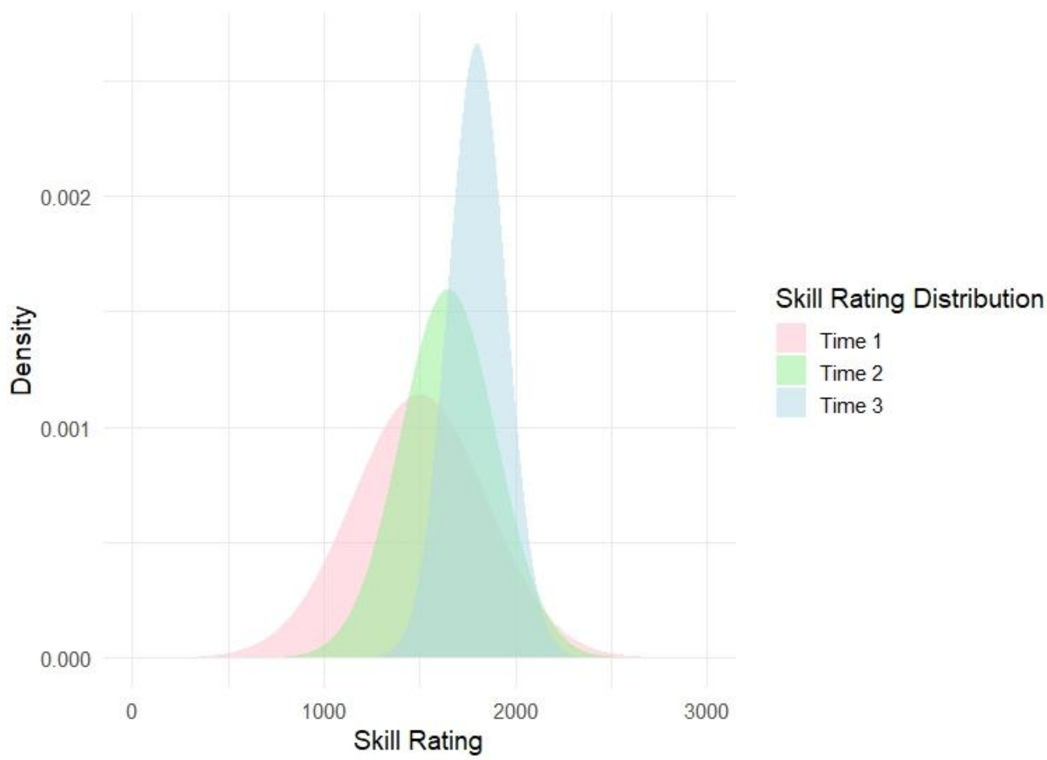

Figure 1. An example competitor skill rating, described as a Gaussian distribution, over three time periods. Note that as time goes on (Time 1 to Time 3), the skill rating distribution is higher (increased skill over time) and narrower (more certainty about the real value of the skill rating). The competitor's real-value skill is somewhere within the distributions at each time interval.

B. Performance as Conditional Probability

Competitors, of course, do not play trading card games in isolation. They play against real opponents. This means that a competitor's *performance,* or $\mathcal{P}$, is based on the probability of Competitor_i's skill *conditional* on Competitor_j's skill, such that:

$$\mathcal{P}_i^k \mid \text{Opponent}_j^k \begin{cases} 1, \text{ if } \text{Skill}_i^k + \epsilon_{i|j}^k > \text{Skill}_j^k \\ 0.5, \text{ if } \text{Skill}_i^k + \epsilon_{i|j}^k \approx \text{Skill}_j^k \\ 0, \text{ if } \text{Skill}_i^k + \epsilon_{i|j}^k < \text{Skill}_j^k \end{cases} \tag{2}$$

Where $\mathcal{P}_i^k$ is Competitor_i's performance conditional on Competitor_j's performance based on their respective Skill and an error term, ε, all in Match *k.* $\mathcal{P}_i^k|\text{Opponent}_j^k$ results in 1 for a "win", 0.5 for a "draw", or 0 for a "loss" for Competitor_i, meaning that Competitor_j, receives the inverse result. The ε error term represents Competitor_i's "luck" (good or bad) conditional on playing Competitor_j in Match *k*. In this study, we do not investigate or control for ε (though this is theoretically accounted for by the fact that skill ratings are distributions and not point estimates).

In Figure 2, we see two sample competitors' skill ratings, call them Competitor_i and Competitor_j, described by Gaussian distributions. Competitor_i's skill is described by $Skill_i \sim \mathcal{N}(1500, 350)$ while Competitor_j's skill is described by $Skill_j \sim \mathcal{N}(1650, 250)$. To solve for the probability of Competitor_i beating Competitor_j in a game match, $P(i|j)$, we will apply Bayes' Rule, as we will see in the following sections.

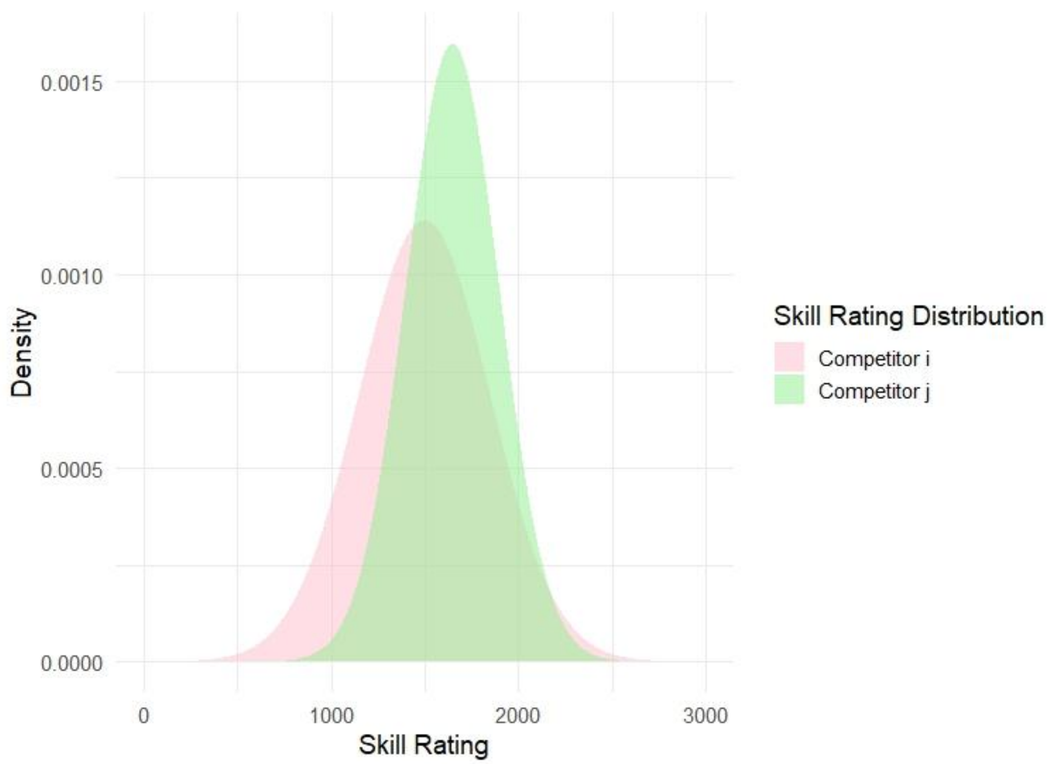

Figure 2. Two competitors' (*i* and *j*) skill rating distributions. In this study, all skill ratings are described by Gaussian distributions.

C. Influence of Deck Selection on Performance

We add to this theory the additional component of deck selection. Competitors in trading card games do not all have the same tools to play against one another like in other games of skill (e.g., chess). Each competitor constructs a deck from available cards and uses this deck against his or her opponents. Deck construction, the contents of each deck, and the randomness involved in shuffling and drawing of cards is what lends trading card games to a higher variability than other games of skill (which further contributes to ε in Equations 2 and 3). Each competitor uses a deck of cards and the selection of deck in a match adds or subtracts from a competitor's total performance. Some decks are simply better than others, or at least, some decks are made to better counter the strategy of others in the competitive landscape.

To account for deck strength, we further elaborate on our Performance Equation, Equation 2, as follows:

$$\mathcal{P}_{i,a}^{k} \mid \mathrm{Opponent}_{j,b}^{k} \begin{cases} 1, \text{ if } \mathrm{Skill}_i^k + \mathrm{Deck}_{a|b}^k + \epsilon_{i|j}^k > \mathrm{Skill}_j^k \\ 0.5, \text{ if } \mathrm{Skill}_i^k + \mathrm{Deck}_{a|b}^k + \epsilon_{i|j}^k \cong \mathrm{Skill}_j^k \\ 0, \text{ if } \mathrm{Skill}_i^k + \mathrm{Deck}_{a|b}^k + \epsilon_{i|j}^k < \mathrm{Skill}_j^k \end{cases} \tag{3}$$

Where Performance is now also conditional on $\mathrm{Competitor}_i$'s Deck_a and $\mathrm{Competitor}_j$'s Deck_b as played in Match *k*.

D. Putting it All Together

We then expect the win probability for $\mathrm{Competitor}_i$ playing against $\mathrm{Competitor}_j$ in Match *k* in which $\mathrm{Competitor}_i$ uses Deck_a and $\mathrm{Competitor}_j$ uses Deck_b to be calculated in the following form:

$$\mathrm{P}\left(\mathrm{Competitor}_{i,a} \middle| \mathrm{Competitor}_{j.b}\right)_k = \frac{\mathrm{Skill}_i^k + \mathrm{Deck}_{a|b}^k + \epsilon_{i|j}^k}{(\mathrm{Skill}_i^k + \mathrm{Deck}_{a|b}^k + \epsilon_{i|j}^k) + (\mathrm{Skill}_j^k)} \tag{4}$$

To embody this competitive theory, we use Mark Glickman's Glicko-2 ratings system [6]. This system gives us the Gaussian competitor skill rating we desire within a robust Bayesian framework. Glicko-2 also provides us with tunable parameters, such as a baseline skill rating and rating deviation for competitor skill, and parameters to control ratings inflation and competitor performance volatility. We will also use the Glicko-2 framework to account for the difference in deck strength between competitors by making a modification to competitor skill ratings conditional on the decks used by either party to a match. A full discussion on our application of the Glicko-2 rating system is given in Section *VII. Competitor Skill Ratings*. Discussion on the issue of deck strength is given in Section *VIII. Deck Strength*.

III. ASSUMPTIONS & TERMINOLOGY

Throughout this paper, we make seven key assumptions and note one special use of terminology:

A. Foreknowledge of Matches and Decks

We assume that we know, in advance, the competitors involved in each match during both the tournament season and during the world tournament. This is a mainstay of sports betting, where events calendars, including the participating teams, are known ahead of time. We also assume that we know, in advance, each competitor's deck type for each match during both the tournament season and during the world tournament. By "in advance," we mean as little as a few minutes beforehand (enough time for the sportsbook's traders to compute win probabilities and set bet prices).

B. "Tournament Season" as a Ratings Period

The "tournament season"—a period lasting twelve consecutive months during which 1,000 games are recorded in each month—will be randomly generated. No bets will be offered or taken during this period. The tournament season exists, in this exercise, to establish a baseline of competitor skill and to establish which of the competitors is qualified to be invited to the world tournament where the model will be fully implemented. In a real-world application, some analysis of prior competitor performance would be necessary to determine competitor skill before bet odds could be computed and offered to customers. For our purposes, the tournament season fulfills this purpose as strictly a ratings period to provide the needed data to extrapolate competitor skill ratings and determine the relative strength of decks.

C. "World Tournament" as Model Experiment

The "world tournament" will exist to demonstrate the TCG sportsbook model's effectiveness. The top performers from the preceding "tournament season" will be invited to participate in this world tournament, during which the amount of bet money placed on each match as well as the game outcomes themselves will be randomly determined.

During the tournament season, "ties" (or "draws") in matches, whereby neither competitor wins or loses, will be recorded, and used to calculate future win probabilities. However, during the world tournament, ties will not be permitted. Each game in the world tournament must end in either a win or a loss for either competitor involved, and thus, for the sportsbook, too, in financial terms. Financial performance from the world tournament will be used to draw inferences about the effectiveness of the model.

D. Wagers on Match Outcomes Only

"Exotic" bets (like parlays, wagers on point totals, "props," etc.) will not be offered. The only bets that our model will allow are wagers on the outcome (win or loss) for a single competitor in each match of the world tournament. The prices of bets will be quoted in the American "moneyline" system [7].

To make a profit, the TCG sportsbook will "overround" the real probabilities it computes for the outcome of a match between competitors. Doing this (along with the "risk management" strategy discussed next) ensures that, give or take, the sportsbook will bake in profit for itself regardless of the outcome of an event. The overround in every case will be 10%. For instance, if a match between Competitor$_i$ and Competitor$_j$ has computed real probabilities of 0.7000 for Competitor$_i$ and 0.3000 for Competitor$_j$, the sportsbook will overround these to 0.7700 for Competitor$_i$ and 0.3300 for Competitor$_i$, for fully implied odds of 1.1 (a mathematical impossibility, given the rules of probability). These overrounded probabilities are then used to price bet odds for customers. Sportsbook customers, in effect, overpay for the odds they are quoted from the house, which is the primary method of profit-making for the sportsbook.

We choose a 10% overround due to the potentially thinly traded nature of a TCG betting market. A cursory review of the implied probabilities from large sportsbooks such as Draft Kings, FanDuel, and BetMGM show an overround of approximately 5% in the most traded markets (e.g., American football, basketball, and baseball). Demand and competition are both high in these markets and thus sportsbook operators can offer lower overround prices to keep their offerings competitive and make money on larger bet volume. A TCG market would be much smaller than these mainstream sports markets; for us to make money with lower volume and higher volatility in outcomes (due to less available data, as acknowledged in *Section IV. Limitations*), we double this overround. A 10% overround provides the TCG sportsbook with a better margin of safety without hugely disrupting the customers' perceived value from potential payouts for winning bets. It is acknowledged here that more work can be done, using real world data, to assess customer sensitivity to the overrounding of bet odds and the subsequent financial impact on the sportsbook.

E. Risk Management Strategy

To mitigate the risk of potential loss to the sportsbook created by unbalanced liability (i.e., too much bettor money on one side of a match), we introduce a simplified risk management strategy.

This simplified risk management strategy helps limit losses that could arise from heavy betting on one outcome, in which the house would stand to make substantial payouts to winning bettors far beyond the money it takes in from all bettors (usually due to heavy betting on an "underdog" for whom the odds are long, and thus, payouts disproportionately large).

As the sportsbook, we can simply limit or refuse action (i.e., bet money from customers) on either side of an event to balance the liability on both sides. In effect, we try to maintain a state of indifference as to the winner (or loser) in any event: our liability is managed and we cannot lose, no matter what, more than our known liability on either side. We keep the system closed, which limits our potential upside if we are right on the outcome of an event, but, more importantly, limits our potential downside if we are wrong.

The mathematical details of this strategy are more fully described in *Section IX. World Tournament*, subsection *B. Simulation of Matches Between Competitors*, subsubsection *Risk Management (Handle: Opponent & Liability: Opponent)* Equations 18 and 19.

The risk management strategy presented in this paper is overly simplified on purpose; the topic of risk management for a sportsbook operator in such a dynamic environment is beyond the scope of this paper and could be investigated in future work.

F. Acceptable Level of Profitability

Profitability will be measured by gross gaming revenue (GGR) margin. GGR is the difference between what the house takes in from bettors from all quoted propositions for an event (the "handle") and what it pays out to bettors who place bets on the winning proposition (the "payout"). For instance, if the sportsbook quotes prices for either Competitor_i or Competitor_j to win a competitive match, the handle will be the total sum of money bettors place on both *i* and *j*. After the match is graded (i.e., a winner determined), the house will pay the bettors of the winning side only. The difference is GGR. The ratio of GGR to the total handle is GGR margin. Thus, we interest ourselves in profitability rated as a percentage of money taken in from bettors; in other words, we are interested in the proportion of how much of this handle we will keep for ourselves as the house.

Our sportsbook's desired level of GGR margin will be 5%. A model that consistently delivers at least 5% GGR margin is a valid model for our purposes; a model that fails to do so is invalid. A full treatment of our inferences about the validity of the model is discussed in *Section XI. Conclusions & Further Research.*

G. Prior Belief

Our Bayesian analysis (see subsection *H. Bayesian Inference,* below) requires us to form a prior belief about the expected GGR margin our model is likely to produce. We form a weakly held prior belief, based on the discussions in subsections *D., E., and F.,* above, that our model should produce a GGR margin of 10% based on our "overrounded" odds of 10%. We will also bake into our prior belief a wide dispersion of potential outcomes for the expected GGR margin. A further discussion on our choices of prior belief will be reviewed in Section *X. Inference,* subsection *A. Prior Distribution.*

H. Bayesian Inference

All inference will be performed using Bayesian reasoning. A Bayesian framework will allow us to draw inferences about the probability distribution of reaching our stated objective and what our model is likely to achieve in that regard if implemented. During the inferential phase, we will fully discuss the setting of a prior belief, the formation of a likelihood distribution based on the observed data from the simulations, and the formation of an updated (posterior) belief about our model's performance.

Due to our choice of prior and an assessment of the observed simulation data, our posterior distribution is not easily expressed in a standard form. We thus employ the Metropolis-Hastings algorithm to sample the posterior distribution from which our inferences are drawn. A full discussion on our inferential process is discussed in Section *X. Inference.*

I. "Player" and "Opponent"

It is also prudent here to discuss the key terminological differences between the terms "player" and "opponent" used throughout this paper and the supporting documentation.

All competitors in this model are referred to as either a "player" or an "opponent." A "player" is the competitor, which is evaluated in any given instance, whether it be for skill, match outcome, or deck use. Each "player," naturally, has an "opponent."

For instance, in given Match k we have two competitors: $Competitor_i$ and $Competitor_j$. If we evaluate $Competitor_i$'s match performance in Match k, $Competitor_i$ is the "player" for that evaluation and $Competitor_j$ is the "opponent." Likewise, if we then evaluate $Competitor_j$'s match performance in Match $k+1$, $Competitor_j$ is referred to as the "player" and $Competitor_i$ as the "opponent." Each "player," thus, has an "opponent," and a given competitor can be either, depending on which side of the evaluation we investigate.

Throughout this paper we use the "player" and "opponent" terminology to refer to competitors on either side of the "ledger," as it were (with the "player" on the left side and the "opponent" on the right side), for any given evaluation of performance, skill rating, match outcome, bet prices, etc. All "players" and "opponents" are "competitors" and "competitors" is used to refer to all game participants irrespective of their place in an evaluation.

IV. LIMITATIONS

In this paper, we acknowledge two key limitations:

A. Legal Feasibility

A sportsbook offering wagers on trading card games would have to pass regulatory muster in any jurisdiction where it wished to offer such wagers to its customers. Trading card games may be especially prone to regulatory scrutiny. Some trading card game participants (that is, the competitors) may be underage. This may cause conflict between regulatory bodies and any sportsbook that wishes to offer its customers wagering opportunities in competitions in which some competitors are not themselves old enough to gamble.

Additionally, even if a gaming jurisdiction permitted wagering on events with underage participants—or all participants were guaranteed to be above some legally required minimum age—regulators would still have to give approval on wagers for this new type of competitive event. In the United States, this would require a state-by-state approval from each state's gaming regulatory authority.

We ignore these issues for the sake of demonstrating the model's technical feasibility and assume that our model operates within a legal jurisdiction that permits its existence. For the sake of the illustration outlined in this paper, all competitors in all games are assumed to be above the legally permitted age to gamble, and thus, likely, admissible as the subjects of bettors' wagering activities. Additionally, we proceed with acknowledgment that a sportsbook offering such wagers to its customers would do so with a full complement of responsible gambling resources and act within responsible gambling norms, like those set out by the American Gaming Association's *Responsible Gaming Code of Conduct* [8].

B. Data Availability

Previously acknowledged is the assumption that all relevant data will be made available to us. This data may not be available, or at least, made available to a sportsbook operator in a timely enough manner given the current state of the industry. For this model to work, the sportsbook must have the following information at hand *before* each match's bet prices can be determined:

1. When will matches be played?
2. Who will play in each match?
3. Which decks will competitors use?

With traditional sports, this type of data is readily available. Game schedules for teams are announced well enough in advance for sportsbooks to determine probabilities and set corresponding bet prices. Realistically, with this TCG model and attendant software systems in place, a sportsbook could compute win probabilities and offer bet prices to customers *within minutes* if not *within seconds,* allowing for near real-time odds pricing for ongoing tournaments. As competitors win or lose matches, win probabilities could be updated and new bet prices offered very quickly.

V. OVERVIEW

To reach the inferential phase, during which we will evaluate the validity of our trading card game sportsbook model, we will use the following methodology:

1. Record all competitor and match data for a single "tournament season" (a twelve-month play period).
2. Use the recorded data from the tournament season to generate competitor skill ratings.
3. Use the recorded data from the tournament season to determine deck strength.
4. Use the recorded data from the tournament season to find and invite the top thirty-two competitors to a "world tournament."
5. Set up a five-round, single elimination world tournament for the thirty-two invitees to play in, during which we:
    a. Use the previously calculated competitor skill ratings to generate win probabilities for each competitor, to include the use of specific deck types, and create a table that lists all pairings of competitors and the win/loss probabilities for each such pairing.
    b. Use the calculated win/loss probabilities to price the odds we offer to sportsbook customers, applying our stated 10% "overround".
    c. Permit bettors to place bets on either side of a match (with random amounts wagered on either side) and aggregate this money as the "handle."
    d. Pay out the winning bettors based on each match outcome (also called the "payout"). Match outcomes are determined randomly.

6. Simulate the outcomes of the five-round world tournament 10,000 times, recording the money wagered, money won/lost by bettors, and money won/lost by the sportsbook for each simulation.

After this is done, we will present the aggregate outcome of these 10,000 world tournament simulations and draw inferences about the validity of our sportsbook model (i.e., whether the model consistently delivers our desired 5% gross gaming revenue margin target). To aid us in this endeavor, we will use Microsoft Excel and the R programming language. Figure 3 gives a high-level overview of data collection and odds making steps we will take before reaching the inferential phase. All files and code used in this paper are available on the project website [4].

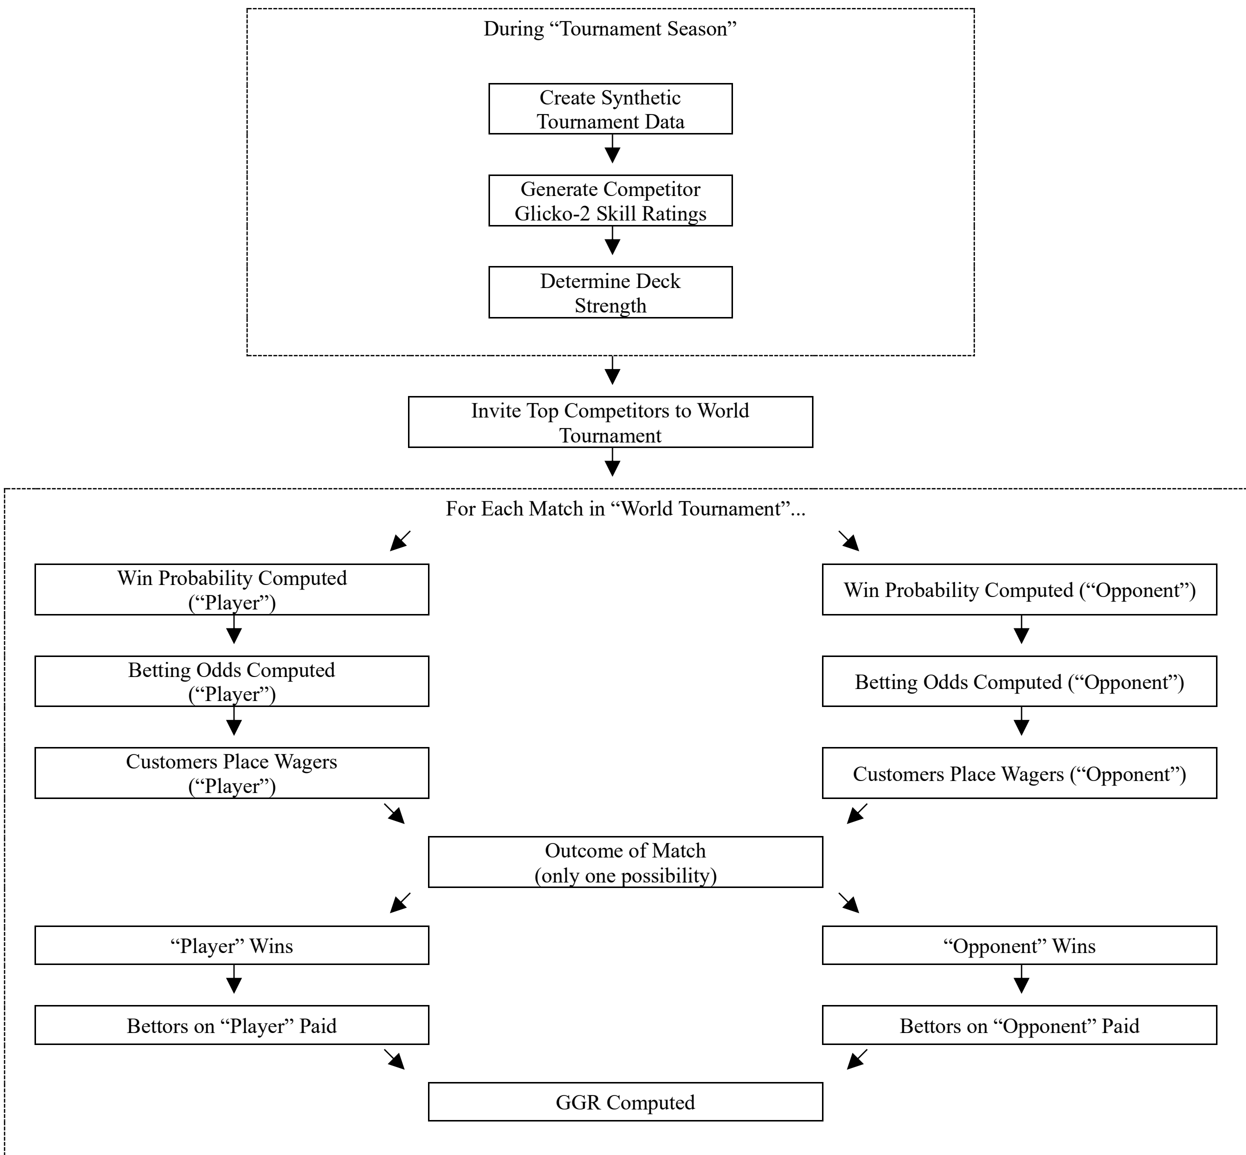

Figure 3. Overview of the methods used in this paper.

VI. TOURNAMENT SEASON

Before we begin to price odds for the world tournament simulations, we must establish competitor skill. To do this, we must first gather sample data on the performance record of the game's competitors (a record of competitor wins, losses, draws, and the opponents they face and decks they use). To do this, we create a synthetic "tournament season."

Our tournament season consists of twelve months ("periods") during a single tournament year. In each of these twelve periods, 1,000 matches between competitors are recorded. In each of these 1,000 matches, two competitors, randomly selected from a pool of five hundred unique competitors ($Competitor_1$,…, $Competitor_{500}$), are pitted against one another (one is the "player", and the other is the "opponent", as discussed previously in Section *III. Assumptions & Terminology,* subsection *I. "Player" and "Opponent"*).

The following two equations are used to determine who plays in each of i matches (where i ranges from 1, …, 1,000) in each of j periods (where j ranges from 1,…, 12). Let R be a discrete random variable over the specified interval in every case:

$$Player_i^j \sim R[1, 500] \tag{5}$$

$$Opponent_i^j \sim R[1, 500], \text{ where } Opponent_i^j \neq Player_i^j \tag{6}$$

Using this method means that not all competitors are recorded as playing in each month of the tournament season, which replicates real trading card game events, in which not all competitors compete in every event. This leaves random (but intentional) gaps in the game record for the tournament season which will create a simulated natural variation in competitors' ratings based, in part, on the frequency of play. Competitors that play more often will have a more reliable skill rating, while those that play less often will have a less reliable skill rating. This, as we will see, is managed easily by the Glicko-2 rating system.

The tournament season data also records which deck each competitor uses ("player deck" or "opponent deck," respectively) during each match. Decks are also assigned numbers, from $Deck_1$, …, $Deck_{15}$, and are assigned to competitors randomly, using the following equation:

$$\mathrm{Deck}_\mathrm{i}^\mathrm{j} = R \in S_j \tag{7}$$

Where R is a randomly assigned discrete value in S_j, the subset of available decks in period j (1,…,12) as shown on Table 1. A changing set of available decks replicates a changing "metagame," in which competitors invent and use new decks over time to solve and defeat previous decks over the course of the tournament season. This makes for an evolving competitive landscape and opens the real-world question (which we address in a later section of this paper) as to how to account for the win/loss probability between decks that have never faced one another in a recorded match.

Table 1. Random Deck Assignment Sets for Equation 7.

Month ("Period")	Available Decks	Subset
1	1, 2, 3, 4	$S_1 \sim R\{1, 2, 3, 4\}$
2	1, 2, 3, 4, 5	$S_2 \sim R\{1, 2, 3, 4, 5\}$
3	2, 3, 4, 5, 6	$S_3 \sim R\{2, 3, 4, 5, 6\}$
4	2, 3, 4, 5, 6, 7	$S_4 \sim R\{2, 3, 4, 5, 6, 7\}$
5	3, 4, 5, 6, 7, 8	$S_5 \sim R\{3, 4, 5, 6, 7, 8\}$
6	3, 4, 5, 6, 7, 8, 9	$S_6 \sim R\{3, 4, 5, 6, 7, 8, 9\}$
7	4, 5, 6, 7, 8, 9, 10	$S_7 \sim R\{4, 5, 6, 7, 8, 9, 10\}$
8	4, 5, 6, 7, 8, 9, 10, 11	$S_8 \sim R\{4, 5, 6, 7, 8, 9, 10, 11\}$
9	5, 6, 7, 8, 9, 10, 11, 12	$S_9 \sim R\{5, 6, 7, 8, 9, 10, 11, 12\}$
10	5, 6, 7, 8, 9, 10, 11, 12, 13	$S_{10} \sim R\{5, 6, 7, 8, 9, 10, 11, 12, 13\}$
11	6, 7, 8, 9, 10, 11, 12, 13, 14	$S_{11} \sim R\{6, 7, 8, 9, 10, 11, 12, 13, 14\}$
12	6, 7, 8, 9, 10, 11, 12, 13, 14, 15	$S_{12} \sim R\{6, 7, 8, 9, 10, 11, 12, 13, 14, 15\}$

To further replicate reality, each deck is also assigned a qualitative "style" typing, which represents the major strategy underpinning the deck's win conditions within the game. This assignment was done randomly, with a uniform probability of each deck, Deck_q, where q is the deck number (Deck_1,…, Deck_{15}):

$$\text{Deck Style}_q \in R(S_d), \text{ where } S_d = \{\text{Aggro, Combo, Control, Mid-Range, Mill}\} \tag{8}$$

Deck styles only factor into our later calculations of determining deck strength against decks that have never faced one another. The summary of deck style types is shown on Table 2.

Lastly, Match k in Period j records a "result," randomly assigned, as either "0", "0.5", or "1", with uniform probability:

$$\text{Game Result}_j^k \in R(S_{GR}), \text{ where } S_{GR} = \{0, 0.5, 1\} \tag{9}$$

A result of "0" means a loss for the "player" and win for the "opponent", a result of 1 means a win for the "player" and a loss for the "opponent", and a result of 0.5 means a tie for both competitors in the match.

The results of the 12,000 tournament season matches, using these random player/opponent and deck assignments, can be found on the website [4a]. A sample of what the data looks like is shown on Table 3.

Table 2. Randomly Assigned Deck Style Types.

Deck	Style	Deck	Style
1	Control	9	Control
2	Combo	10	Control
3	Aggro	11	Mill
4	Combo	12	Mill
5	Mid-Range	13	Combo
6	Mid-Range	14	Mid-Range
7	Control	15	Aggro
8	Combo	16*	Control

*$Deck_{16}$ is not seen until the world tournament, after the tournament season. The inclusion of this deck's style type assignment is included on this table for convenience. Handling of $Deck_{16}$ is discussed in Section *VIII. Deck Strength*.

Table 3. Sample of Tournament Season Data.

Match	Period	Player	P_Deck	PD_Style	Opponent	O_Deck	OP_Style	Result
1	1	121	2	Combo	444	1	Control	0
2	1	448	1	Control	495	4	Combo	0
3	1	461	4	Combo	11	2	Combo	0
4	1	356	1	Control	312	1	Control	0.5
5	1	8	4	Combo	309	3	Aggro	0
6	1	80	2	Combo	125	4	Combo	0
7	1	483	2	Combo	266	2	Combo	0
8	1	355	3	Aggro	228	1	Control	1
9	1	331	1	Control	412	2	Combo	1
10	1	452	3	Aggro	256	3	Aggro	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

VII. COMPETITOR SKILL RATINGS

After the tournament season data have been generated, we move to establishing competitor skill ratings within the Glicko-2 framework discussed in *Section II. Theory*. The calculation of competitors' Glicko-2 ratings was performed in `R` with the `PlayerRatings` package [9].

To generate the Glicko-2 ratings in `R`, the parameters shown on Table 4 were chosen to establish a baseline for all competitors at the outset of the tournament season. We allow the outcomes of the recorded matches in the tournament season to change the ratings, deviations, and volatility per the Glicko-2 system as matches take place. The constant, c, is left unchanged in all calculations throughout the model's iteration.

Thus, all competitors begin the tournament season on even footing: a skill rating of 1500, a rating deviation of 350, and a volatility of 0.6. These values fall in line with Glickman's own recommendations for the system [6]. As wins, losses, and ties are recorded as the tournament season progresses, these values change for each of the five hundred competitors based on each competitor's individual performance. (As mentioned previously, the rating inflation constant, *c,* is held the same throughout the model.)

Table 4. Baseline Glicko-2 System Variable Values.

Variable	Value
Skill Rating	1500
Rating Deviation	350
c (rating inflation constant)	60
τ (rating volatility)	0.6

A sample of the skill ratings generated by the `PlayerRatings` package and the Glicko-2 system can be seen on Table 5. The `R` code used to do this can be seen on the website along with the "final standings" of all the competitors at the end of the tournament season [4b] [4c]. These scores will be used to compute win probabilities for competitors invited to the world tournament.

Table 5. Sample of Final Tournament Season Competitor Glicko-2 Ratings (from highest to lowest).

Player	Rating	Deviation	Volatility	Games	Win	Loss	Draw	Lag
42	1826.75	182.7026	0.6167	28	18	8	2	0
67	1797.878	136.0028	0.5779	50	26	15	9	0
37	1771.898	132.1543	0.5946	49	28	17	4	0
407	1744.987	132.9164	0.5879	44	24	15	5	0
284	1731.62	152.4804	0.5781	36	17	14	5	0
287	1723.616	119.6335	0.6019	49	28	15	6	0
453	1721.794	128.9282	0.5714	49	28	15	6	0
215	1714.139	144.3322	0.5644	49	25	17	7	0
370	1703.363	147.9625	0.5681	38	18	13	7	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

VIII. DECK STRENGTH

With competitor skill firmly established by the Glicko-2 framework, we now turn to determine deck strength. By this we mean to answer the question *"how much does a competitor's choice of deck contribute (or detract) from his or her total 'performance' in a match?"*

A. Adjusting Glicko-2 Ratings

The Glicko-2 system already provides us with a way of calculating win probabilities between competitors, conditional on their playing against one another. To account for a competitor's deck contribution to his or her performance, we seek to adjust the ratings of the competitors in a match so that this adjustment reflects the decks' effect on the match outcome.

To determine how much adjustment to competitor ratings should be made, we consider the following: two competitors, call them Competitor_i and Competitor_j, face one another in a match. Both competitors have *identical* skill ratings and ratings deviations. The win probability of Competitor_i against Competitor_j is 0.5000, that is, $P(i|j) = 0.5000$. The only difference between these two competitors will be the decks that they use. Suppose $Competitor_i$ uses $Deck_a$ and $Competitor_j$ uses $Deck_b$. Also suppose that, given our records, $Deck_a$ beats $Deck_b$ with 0.6100 probability, that is, *P(a|b) = 0.6100.* This then holds that $Competitor_i$ using $Deck_a$ against $Competitor_j$ using $Deck_b$ has a win probability of 0.6100, that is, *P(i,a | j,b) = 0.6100.* Again, with two dead even skill ratings and rating deviations, the only decisive factor influencing win probability is the win probability *between the two decks in use.*

In a match such as Match 1 on Table 6, we do not offset either competitor's Glicko-2 skill rating: the matchups are dead even. However, as shown in Matches 2 through 9 on Table 6, different deck matchups result in offsetting the competitors' skill ratings by a certain amount.

Table 6. Nine Sample Matches Between Competitors and the Skill Rating Offset for Each.

Match	$Skill_i$, RD_i	$Skill_j$, RD_j	P(a|b)	Γ	P(i|j)	P(i,a|j,b)
1	1500, 350	1500, 350	0.50	0	0.5000	0.5000
2	1500, 350	2000, 350	0.50	0	0.1272	0.1272
3	2000, 350	1500, 350	0.50	0	0.8728	0.8728
4	1500, 350	1500, 350	0.61	130	0.5000	0.6226
5	1500, 350	2000, 350	0.61	130	0.1272	0.1939
6	2000, 350	1500, 350	0.61	130	0.8728	0.9188
7	1500, 350	1500, 350	0.39	-130	0.5000	0.3774
8	1500, 350	2000, 350	0.39	-130	0.1272	0.8012
9	2000, 350	1500, 350	0.39	-130	0.8728	0.8061

On Table 6, we list the skill rating and rating deviation (RD) for each competitor, *i* and *j,* as well as the win probability of $Deck_a$ against $Deck_b$, *P(a|b),* the unadjusted win probability of $Competitor_i$ against $Competitor_j$ without taking decks into account, *P(i|j),* and the adjusted win probability of $Competitor_i$ using $Deck_a$ against $Competitor_j$ using $Deck_b$, *P(i,a|j,b).* The RD for both competitors in each sample match is kept at 350 (the maximum allowable by the Glicko-2 system), which represents for us in these examples an environment of maximum uncertainty.

The term "Γ" on Table 6 is the adjustment term. This is the amount that $Competitor_i$'s skill rating is adjusted to reflect the probability of $Deck_a$ against $Deck_b$ in a match against $Competitor_j$. Thus, a positive Γ adds to $Competitor_i$'s Glicko-2 skill rating, and consequently, $Competitor_i$'s win probability, and a negative Γ subtracts from $Competitor_i$'s skill rating, and consequently subtracts from his or her win probability. The Γ term is temporary and only used to adjust for the influence of deck strength on a conditional matchup of two competitors in each match.

The amount of offset, in each case, is described by the following equation:

$$\Gamma \cong 290.534 \times ln\left(\frac{P(a|b)}{1-P(a|b)}\right) \tag{10}$$

Using Equation 10 to calculate the Γ for a deck with a 0.6100 win probability, that is, *P(a|b) = 0.6100,* we find a Γ of approximately 130. Conversely, to calculate the Γ for a deck with a 0.3900 win probability, that is, *P(b|a) = 0.3900,* we find a Γ of approximately -130. This is the amount $Competitor_i$'s Glicko-2 skill rating is adjusted in a match given the use of these decks (either +130 if using $Deck_a$ against $Deck_b$ or -130 is using $Deck_b$ against $Deck_a$). These examples are illustrated in Table 6.

To visualize the impact of Γ on competitors' win probabilities, we set two evenly matched competitors against one another in a series of matches, each with a different Γ ranging from -1500 to 1500, as shown in Figure 4. The result is a logistic curve that expresses the marginal change in win probability for each adjustment increment of Γ. This computation was performed in `R` [4d].

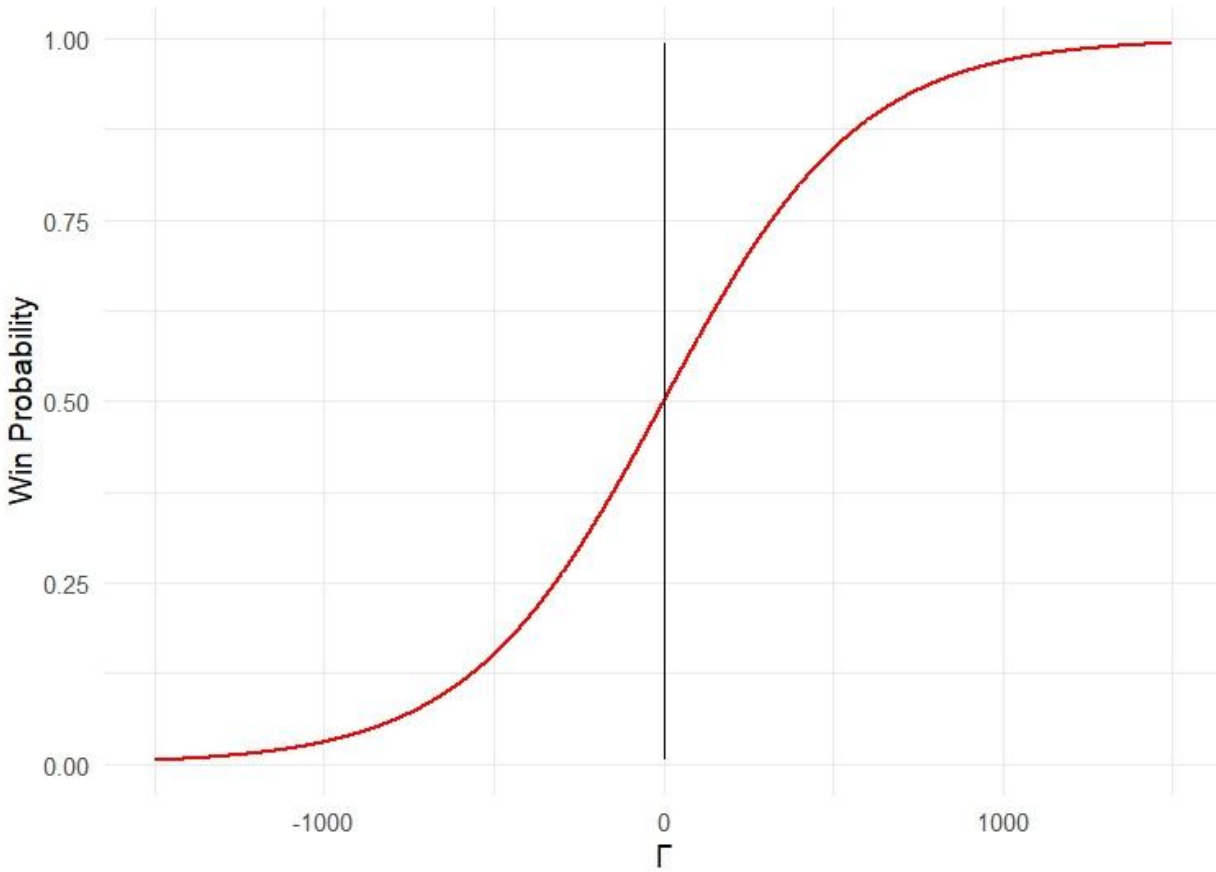

Figure 4. The effect of deck strength, or "Γ," on a competitor's win probability given an even matchup.

The sharpest changes in win probability happen at or near $\Gamma = 0$, when decks are evenly matched. This means that, theoretically, at or around $P(i|j) = 0.5000$, $P(a|b)$ is decisive, because the competitors are otherwise evenly matched. At the upper or lower bounds of the curve, changes in $P(a|b)$ are less decisive. This makes intuitive sense because skill should be the controlling factor. At very low levels of skill, the choice of decks should mean little as the competitor isn't skilled enough to make full use of the available game tools, while at higher levels of skill, a competitor should perform well relative to weaker competitors no matter the choice of deck (in other words, that the choice of deck accounts for increasingly less of the higher-skilled competitor's performance, the larger the gap in skill).

Consider the best competitor in a game, called $Competitor_m$, using the worst deck in the game, called $Deck_x$, in match against the worst competitor in the game, called $Competitor_n$, using the best deck in the game, called $Deck_y$. The skill gap is so great that *m* has *never* lost a match and has a Glicko-2 rating of 3,000 and *n* just learned the rules of the game five minutes before sitting down to play. How much does the deck of either competitor matter in this case?

Consider the four scenarios described in Table 7.

Table 7. Illustration of Deck Strength for Best and Worst Competitors in a Game.

Scenario	**State**	**Γ**	**Probability**	**Change***
1	P(n,x|m,y)	-638	0.0003	-99.14%
2	P(n,y|m,x)	+638	0.0349	+11,533%
3	P(m,x|n,y)	-638	0.9651	-3.46%
4	P(m,y|n,x)	+638	0.9997	+3.59%

*The "Change" measures the percentage change in probability between corresponding scenarios. Thus, the Change shown on this table compares the change in probability between Scenarios 1 and 2 and between Scenarios 3 and 4, respectively.

In each case illustrated by Table 7, $Competitor_m$ has a skill rating of 3,000 and rating deviation (RD) of 350 and $Competitor_n$ has a skill rating of 1,500 and RD of 350. The deck matchup is so lopsided that $P(x|y)$ is just 0.1000 (and conversely, $P(y|x)$ is 0.9000). This means that, ceteris paribus, $Deck_x$ beats $Deck_y$ with only 0.1000 probability.

The scenarios in Table 7 clearly illustrate that skill is the controlling factor. While the choice of deck increases *n's* win probability over eleven times from *x* to *y*, even this much help from deck choice does little when compared to *m's* superior skill. Likewise, even if *m* has the worst deck in the game, he or she has the overwhelming probability of winning. Only at close levels of skill rating between competitors, as previously illustrated in the examples in Table 6 and the logistic curve described in Figure 4, should deck choice be the decisive factor. Deck choice does have an influence, and its effects are on a gradient somewhere between the extreme examples of Scenario 1 and Scenario 4 on Table 7.

B. Computing Win Probabilities for Known Decks

With our offset model determined, we must now also know how to compute deck win probabilities. We perform a piecewise comparison by matching up each deck ($Deck_1$, …, $Deck_{15}$) against each other deck used during the tournament season and compute the win percentage.

$$D_{a|b} = \frac{D_a}{D_a + D_b} \tag{11}$$

Where $D_{a|b}$ is the win probability of $Deck_a$ conditional on $Deck_b$. This is done in Microsoft Excel for all fifteen decks used in the tournament season, and the results of the calculations are shown on Table 8. In a real-world application, this piecewise comparison could be updated automatically using software to update $D_{a|b}$ for each deck matchup as game matches concluded and data made available.

Table 8. Deck Win Probabilities from the Tournament Season.

Deck	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	0.5000	0.5618	0.5434	0.5050	0.6029	-	-	-	-	-	-	-	-	-	-
2	0.4382	0.5000	0.5047	0.4969	0.5288	0.5271	0.4783	-	-	-	-	-	-	-	-
3	0.4566	0.4953	0.5000	0.5326	0.4907	0.4961	0.5216	0.5245	0.5357	-	-	-	-	-	-
4	0.4950	0.5031	0.4674	0.5000	0.4736	0.5053	0.4793	0.4806	0.4330	0.5063	0.4286	-	-	-	-
5	0.3971	0.4712	0.5093	0.5264	0.5000	0.5111	0.4547	0.4832	0.5253	0.4760	0.5595	0.4821	0.6136	-	-
6	-	0.4729	0.5039	0.4947	0.4889	0.5000	0.5095	0.5241	0.5433	0.5085	0.5224	0.5045	0.5137	0.4286	0.7105
7	-	0.5217	0.4784	0.5207	0.5453	0.4905	0.5000	0.5116	0.4036	0.5452	0.5891	0.4494	0.4519	0.5306	0.4773
8	-	-	0.4755	0.5194	0.5168	0.4759	0.4884	0.5000	0.5503	0.5444	0.4795	0.5804	0.5667	0.3684	0.6389
9	-	-	0.4643	0.5670	0.4747	0.4567	0.5964	0.4497	0.5000	0.5402	0.5039	0.4491	0.4865	0.5439	0.7292
10	-	-	-	0.4938	0.5240	0.4915	0.4548	0.4556	0.4598	0.5000	0.5000	0.4592	0.5565	0.5000	0.4412
11	-	-	-	0.5714	0.4405	0.4776	0.4109	0.5205	0.4961	0.5000	0.5000	0.4745	0.4722	0.3909	0.6250
12	-	-	-	-	0.5179	0.4955	0.5506	0.4196	0.5509	0.5408	0.5255	0.5000	0.5083	0.6538	0.4412
13	-	-	-	-	0.3864	0.4863	0.5481	0.4333	0.5135	0.4435	0.5278	0.4917	0.5000	0.4649	0.3235
14	-	-	-	-	-	0.5714	0.4694	0.6316	0.4561	0.5000	0.6091	0.3462	0.5351	0.5000	0.5833
15	-	-	-	-	-	0.2895	0.5227	0.3611	0.2708	0.5588	0.3750	0.5588	0.6765	0.4167	0.5000

C. Computing Win Probabilities for Unknown Decks

After deck strength has been determined for *previously played* decks, we turn now to how to decide what win probability to assign a deck that has *never yet been played in a tournament setting*. In our world tournament simulations, we will introduce a new deck ($Deck_{16}$) that has never been used by a competitor before in a recorded match. This is a common occurrence in real trading card games, whereby competitors bring new strategies to events to counter existing strategies.

If a new, unseen, deck enters the metagame, we will assign a win probability respective to the two decks' style types, based on the two competing styles' win probabilities. Let $T_{a|b}$ be the win probability of $Deck_a$ conditional on $Deck_b$, using the historical data for style types:

$$T_{a|b} = \frac{T_a}{T_a + T_b} \tag{12}$$

This calculation is performed for all style types, as shown on Table 8.

These probabilities will be used for any matchup in the world tournament between one or more decks facing one another for which no record exists in the tournament season data. (In other words, Equation 12 and Table 9 apply to any match that involves the yet-to-be-tested $Deck_{16}$). In a real-world application, these calculations could also be performed dynamically by software based on updated information.

Table 9. Win Probabilities for Deck Style Types.

Deck Style	Aggro	Combo	Control	Mid-Range	Mill
Aggro	0.5000	0.5182	0.4838	0.4845	0.4595
Combo	0.4818	0.5000	0.4938	0.4931	0.5118
Control	0.5162	0.5062	0.5000	0.5052	0.4978
Mid-Range	0.5155	0.5069	0.4948	0.5000	0.5112
Mill	0.5405	0.4882	0.5022	0.4888	0.5000

With competitor ratings and deck strength solved, we move on to setting up a world tournament for the top performers from the preceding tournament season.

IX. WORLD TOURNAMENT

A. Top Competitors as Invitees to World Tournament

To set up the world tournament simulations, we must first determine which of the competitors in the preceding tournament season were the top performers. For the purposes of the tournament organizer, "performance" is a function of the following:

1. Highest number of wins; then
2. Lowest number of losses; then
3. Lowest number of draws.

The tournament organizer (theoretically, the producer of the trading card game itself) does not care about (nor probably has the capacity or inclination to figure) performance scores such as

Glicko-2 ratings. Trading card games invite competitors to large tournaments, like world tournaments, based on a win/loss record during a tournament season. Our simplified method gives us the top thirty-two competitors, which are summarized in Table 10. Each of the competitors is "seeded," meaning that the best performer is "Seed 1", the second best performer is "Seed 2", etc. In each of our simulations for the world tournament, Seed 1 will face Seed 32, Seed 2 will face Seed 31, etc., in the first round. The winners from each of these initial matchups will move on to successive rounds, as shown in Figure 6.

B. Random Deck Assignment

Additionally, a random deck is assigned to each competitor at the outset of all simulations, such that:

$$\text{Deck}_i^w \in R(S_w), \text{ where } S_w = \{11,12,13,14,15,16\} \tag{13}$$

Where Deck_i^w is the random deck assignment for each competitor in the world tournament *(w)* for the i^{th} seeded competitor (i = 1, …, 32). This random deck assignment will hold constant throughout all matches in all 10,000 simulations, as shown on Table 10.

The set S_w in Equation 13 contains Deck_{16}. This means that we will see Deck_{16} appear in the world tournament for the first time. Any match in the world tournament involving any decks other than Deck_{16} (that is, Deck_1, …, Deck_{15}) will calculate the deck strength Γ variable using Equation 11 (and Equation 10) while any matchup involving Deck_{16} (for one or both sides of the match) will calculate the deck strength Γ variable using Equation 12 (and Equation 10), as previously discussed and as illustrated in Figure 5.

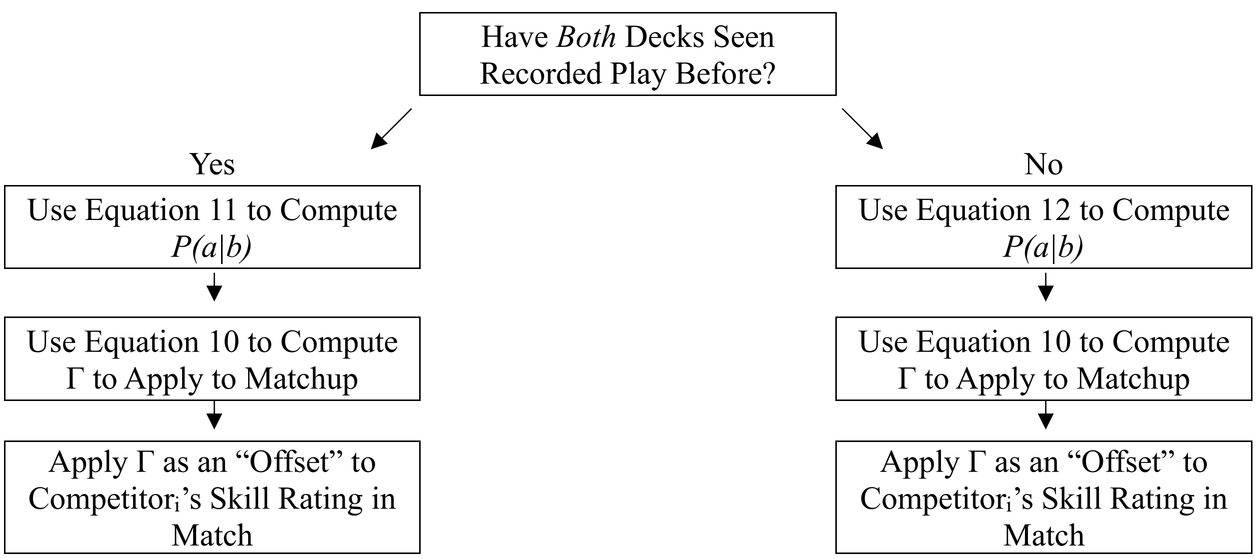

Figure 5. Flow of how to compute Γ for a matchup between Competitors *i* and *j* based on whether both decks (*a* and *b*) have been seen in rated tournament play before.

Figure 6. "Seeded" bracket for the world tournament showing five rounds of single elimination matches. The pairings for the first round in each of the 10,000 simulations are identical to those shown here. Competitors that move on to subsequent rounds vary by simulation, based on the computed win probabilities for each competitor and the random number generation used to determine the winner of each match.

Table 10. Top Thirty-Two "Seeded" Competitors Invited to the World Tournament.

Seed	Player #	Deck #	Seed	Player #	Deck #
1	216	12	17	355	14
2	12	13	18	50	15
3	253	16	19	269	15
4	281	12	20	496	12
5	115	13	21	90	16
6	353	16	22	185	13
7	455	12	23	402	11
8	320	12	24	134	13
9	429	13	25	490	16
10	104	14	26	181	15
11	344	16	27	161	13
12	444	13	28	287	12
13	7	14	29	453	13
14	133	11	30	136	14
15	478	11	31	37	11
16	363	11	32	71	13

C. Competitor Win Probability Matrix

Using `R`, we take the competitors' Glicko-2 scores and compute their win probabilities against one another. We construct a 32x32 matrix giving the win probability of each competitor on the y-axis against each other competitor on the x-axis (including a competitor facing himself or herself), using the tournament season skill ratings and computed Γ values reflecting each competitor's deck assignment. The `R` code to perform this operation, the Γ values computed to apply to win probabilities, and the full resulting matrix can be found on the website [4e-h]. A summary of this matrix can be found on Table 11.

Table 11. Sample of World Tournament Competitor Win Probability Matrix.

	Seed	**1**	**2**	**3**	**4**	**5**	**…**	**32**
Seed	**Competitor**	**216**	**12**	**253**	**281**	**115**	**…**	**71**
1	**216**	0.5000	0.5015	0.5175	0.6842	0.5153	…	0.6562
2	**12**	0.4986	0.5000	0.4582	0.6861	0.5141	…	0.6576
3	**253**	0.4826	0.5418	0.5000	0.6696	0.5555	…	0.6917
4	**281**	0.3158	0.3140	0.3305	0.5000	0.3262	…	0.4701
5	**115**	0.4848	0.4859	0.4445	0.6739	0.5000	…	0.6450
6	**353**	0.3472	0.3999	0.3625	0.5339	0.4133	…	0.5604
7	**455**	0.4299	0.4301	0.4473	0.6249	0.4441	…	0.5945
8	**320**	0.3977	0.3973	0.4141	0.5876	0.4106		0.5577
⋮	**⋮**	⋮	⋮	⋮	⋮	⋮	…	⋮
32	**71**	0.3439	0.3424	0.3083	0.5301	0.3550	…	0.5000

These probabilities will be used in the simulations to determine the fair probabilities of the outcomes of each simulated match.

D. Simulation of Matches Between Competitors

With the thirty-two invitees chosen, the win probabilities for each competitor matchup determined (including factors for deck strength), and the tournament bracket properly seeded, we work next to construct the 10,000 simulations of the world tournament. Each of these simulations, as outlined previously and illustrated in Figure 6, will contain thirty-one matches to decide a winner. With 10,000 world tournament simulations, we will end up with 310,000 unique matches in our final dataset.

For each match in each simulation, the following twenty-two variables are presented.

i. Simulation

The numbered identity of the simulation (from 1, …, 10000).

ii. Round

Within each of the 10,000 simulations, there are five rounds (1, …, 5) representing matches between competitors and one dummy round (Round 6 in each simulation). The pairings follow Figure 6 for each simulation. The dummy round acts as a placeholder to show the winner of the tournament in each simulation, as this competitor is the winner of the fifth and final round of each simulated tournament. (Thus, in each dummy round, both competitors are the same: this is the winner of each final match, round five, in each simulation.)

This leaves us with *k* Matches in the full dataset.

iii. Match

The numbered game Match (1,…, 310000) in the world tournament simulation dataset. Note that in any "dummy" round, this variable is denoted by D_1, …, D_{10000} (as each simulation has its own dummy round listing the winner of that simulated world tournament).

"Dummy" rounds are not used in any analysis in this paper. They are merely placeholders.

iv. Player

The competitor (by number) representing the "player."

v. Player Win Probability

The "player's" real win probability conditional on the "opponent," taken from Table 11. This variable is denoted by $P(p_k)$.

vi. Overround Player Probability

The "over rounded" (or overpriced) win probability set by the sportsbook. This is simply 110% the Player Win Probability, to represent a 10% "overround". This variable is denoted by $P(p_k)'$.

vii. Moneyline Operator: Player

This is a symbol that is either a "+" (in the case of the Player Win Probability >0.5000, to indicate a "favorite" in a match) or a "-" (in the case of the Player Win Probability <0.5000, to indicate an "underdog"). These operators assist in understanding the pricing of odds in the American moneyline system [7].

viii. Moneyline: Player

The moneyline odds for bets placed on the "player" in Match *k*. This is the sportsbook quoted price to bettors for a wager on the "player" to win the match.

This is calculated as:

$$M_p^k \begin{cases} \dfrac{100}{P(p_k)'} - 100, \text{ if } P(p_k) < 0.5 \\ \dfrac{100 \times P(p_k)'}{1 - P(p_k)'}, \text{ if } P(p_k) > 0.5 \end{cases} \tag{14}$$

Where $P(p_k)$ is the real player win probability and $P(p_k)'$ is the overround player probability in Match *k*.

ix. Handle: Player

This is the amount of bet money from customers placed on the "player" in the match. This is a random dollar amount between $5,000 and $10,000, such that:

$$H_p^k = R(5000, 10000) \tag{15}$$

Where H_p^k is the handle on the "player" in Match *k*.

x. Liability: Player

This is the liability owed by the house to customers if the "player" wins the match, such that:

$$L_p^k \begin{cases} H_p^k + \left(H_p^k \times \dfrac{M_p^k}{100}\right), \text{ if } P(o_k) < 0.5 \\ \dfrac{H_p^k \times (100 + M_p^k)}{M_p^k}, \text{ if } P(o_k) > 0.5 \end{cases} \tag{16}$$

Where L_p^k is the sportsbook's liability on the "player," M_p^k is the moneyline odds price on the "player," H_p^k is the total handle on the "player," and $P(p_k)$ is the "player" real win probability, all in Match *k*.

xi. Through xvii.

The "opponent" has his or her own win probability, overround probability, moneyline operator, moneyline odds, handle, and liability. These function identically to those of the "player," but for the competitor that represents the "opponent" in each match.

Note that the opponent win probability in Match k, $P(o_k)$ is simply:

$$P(o_k) = 1 - P(p_k) \tag{17}$$

The variables $P(o_k)'$, M_o^k, H_o^k, and L_o^k, replace $P(p_k)'$, M_p^k, H_p^k, and L_p^k, respectively in Match k.

Risk Management (Handle: Opponent & Liability: Opponent)

Note that, as previously mentioned, we have a "risk management" strategy baked into our model. To reflect this strategy, we first look at the "opponent" liability variable, L_o^k:

$$L_o^k = L_p^k \times \left(1 + \frac{R(-0.25, 0.25)}{100}\right) \tag{18}$$

This means that we take the "player liability," L_p^k, in Match k and simply apply to it a random factor between -25 to +25%. This reflects our traders' ability to balance liability on both sides of a match, keeping the house indifferent to the outcome.

Next, we work backward to compute the "opponent" handle, H_o^k, thusly:

$$H_o^k \begin{cases} \dfrac{100 \times L_o^k}{M_o^k + 100}, \text{if } P(o)_k < 0.5 \\ \dfrac{L_o^k \times M_o^k}{100 + M_o^k}, \text{if } P(o)_k > 0.5 \end{cases} \tag{19}$$

xviii. Random Number

Here, a random number (R_1^k) is generated between 0 and 1 to help determine the result of the Match k between competitors, such that:

$$R_1^k = R(0,1) \tag{20}$$

Where R is a random discrete value on that interval.

ixx. Result

This variable (R_2^k) records the result of Match k, as determined by the random number generated from variable R_1^k.

If the random number is < "player's" win probability, a "1" is recorded (for a "player" win), if the random number is > "player's" win probability, a "0" is recorded (for an "opponent" win), such that:

$$R_2^k = \begin{cases} 0, \text{if } R_1^k > P(p_k) \\ 1, \text{if } R_1^k < P(p_k) \end{cases} \tag{21}$$

No draws are permitted. A competitor must win each match conclusively.

xx. Handle Total

The handle total *($H_{t,}^{k}$)* is the sum of variables H_p^k ("player" handle) and H_o^k ("opponent" handle) in Match *k*. This is the total amount of bet money taken in by the house from bettors on both sides of each match.

$$H_t^k = H_p^k + H_o^k \tag{22}$$

xxi. Payout

The payout (ϕ_k) is the amount to be paid by the house to the winning customers that placed a bet on the correct result of Match *k*, such that:

$$\phi_k = \begin{cases} L_o^k, \text{ if } R_2^k = 0 \\ L_p^k, \text{ if } R_2^k = 1 \end{cases} \tag{23}$$

xxii. GGR

"Gross gaming revenue" (Σ_k) is the difference between H_t^k (handle total) and ϕ_k (payout) in Match *k*, such that:

$$\Sigma_k = H_t^k - \phi_k \tag{24}$$

A positive value for Σ_k represents a financial gain for the house while a negative value for Σ_k represents a financial loss for the house on Match *k*.

D. Putting It All Together

The calculations using these variables are performed in Excel and the full simulation workbook is available on the website [4i]. A sample of the simulation data is shown in Table 12. A single draw of 10,000 simulations, representing 310,000 individual matches and their outcomes is also available on the website [4j-k]. This draw of 310,000 simulated matches is used throughout the following sections as our model data.

With the simulations drawn, we move next to our inferential phase.

Table 12. Sample Data from the 10,000 Simulations.

Simulation	Round	Match	Player	P_Win_Prob	OR_P_P	MLOP	Moneyline_P	Handle_P	Liability_P	…	Rand_Num	Result	Handle_T	Payout	GGR
1	1	1	216	0.6562	0.7218	-	259.51	5480	7591.67	…	0.2306	1	8149.87	7591.67	558.21
1	1	2	363	0.4956	0.5451	+	83.44	7183	13176.71	…	0.2638	1	13470.79	13176.71	294.08
1	1	3	429	0.5515	0.6067	-	154.23	5135	8464.50	…	0.9917	0	9686.78	9226.30	460.48
1	1	4	320	0.4687	0.5156	+	93.95	7293	14144.94	…	0.9385	0	15063.45	13296.24	1767.20
1	1	5	281	0.1872	0.2059	+	385.75	5480	26619.22	…	0.0042	1	34755.40	26619.22	8136.18
1	1	6	7	0.3906	0.4296	+	132.75	5409	12589.61	…	0.256	1	14692.53	12589.61	2102.92
1	1	7	444	0.5749	0.6324	-	172.05	5806	9180.66	…	0.4468	1	10613.85	9180.66	1433.19
1	1	8	115	0.3041	0.3345	+	198.94	9262	27687.66	…	0.3502	0	31092.25	28518.29	2573.97
1	1	9	12	0.3112	0.3423	+	192.16	8078	23600.64	…	0.9707	0	23457.12	20296.55	3160.57
1	1	10	478	0.7013	0.7714	-	337.54	8602	11150.46	…	0.3164	1	12412.04	11150.46	1261.59
1	1	11	104	0.8342	0.9176	-	1113.44	8997	9805.04	…	0.688	1	10964.40	9805.04	1159.36
1	1	12	455	0.2957	0.3253	+	207.45	5135	15787.37	…	0.5616	0	18222.29	16892.49	1329.80
1	1	13	253	0.5117	0.5628	-	128.75	9769	17356.41	…	0.8471	0	16947.75	13364.44	3583.31
1	1	14	133	0.6362	0.6998	-	233.13	9080	12974.90	…	0.7292	0	14428.16	13364.15	1064.01
1	1	15	344	0.5338	0.5871	-	142.21	8783	14959.09	…	0.765	0	15380.92	12864.82	2516.10
1	1	16	353	0.3396	0.3735	+	167.72	8442	22601.05	…	0.2724	1	23711.82	22601.05	1110.76
1	2	17	216	0.4255	0.4680	+	113.68	5442	11628.24	…	0.4602	0	13746.45	13139.91	606.53
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	…	⋮	⋮	⋮	⋮	⋮
10000	5	310000	216	0.5550	0.6105	-	156.76	8249	13511.09	…	0.9884	0	15523.52	14862.20	661.32

X. INFERENCE

With our data collected from the 10,000 world tournament simulations, we turn now to perform our Bayesian inference. We will first form a prior belief about the efficacy (i.e., profitability) of the sportsbook model, then evaluate the data and form a likelihood, and finally, reconcile these two (prior and likelihood) to form a conclusion (posterior belief) about what kind of financial performance we can expect from our model. We look to understand the profitability of our model as measured by gross gaming revenue (GGR) margin, the dispersion of GGR margin, and the probability that we meet our GGR margin target (>5%).

It should be noted that our prior belief will take the form of a Gaussian distribution based on the logic of our model. The data itself, as we will see, has a wide dispersion with many outliers. This leads us to form a likelihood in the form of a Cauchy distribution. Thus, our posterior distribution will not have a closed-form solution. We will use the Metropolis-Hastings (MH) algorithm and a Markov chain Monte Carlo (MCMC) strategy to sample from this open-form posterior distribution to arrive at our inferences about the efficacy of the model.

It is tempting to fit our beliefs into purely Gaussian distributions, which are more convenient to work with and easier to compute. Our "risk management" strategy, in theory, should protect the house against surprise financial outcomes due heavy bettor money on one side of an event, creating massive liability. Despite this, we take a more conservative approach by allowing for more extreme potential results. We trust our traders to apply our risk management strategy well, but we will help safeguard against "overfitting" our modelled assumptions by forming an open-form posterior which has "fatter" tails than a Gaussian distribution, and thus, allows for the higher probability of extreme outcomes.

A. Prior Distribution

For our prior belief, we choose a Gaussian distribution based on our theoretical "overround." Our prior takes the following form:

$$p(\theta) \sim N(0.1000, 0.2000)$$

The mean (μ) for our prior is 0.1000, our overround for priced odds, and the standard deviation (σ) is 0.2000, representing a wide dispersion of potential outcomes. This means that we expect most of our GGR margin to fall somewhere between -30 and +50%. This constitutes a weakly informative prior about the potential for GGR margin produced by our model.

Our prior belief can be visualized by the distribution illustrated in Figure 7.

Figure 7. Prior (Gaussian) distribution for potential GGR margins based on the theoretical 10% overround.

B. Likelihood Distribution

We take caution in assessing the data gathered from the simulations [4k]. It would be tempting to use the simulated GGR results from the 10,000 world tournament simulations and not the 310,000 individual simulated matches. Using the 10,000 world tournament simulations means that each of these 10,000 simulations themselves are *averages of 31 matches each.* This skews the reported results making them look more favorable. We thus assess the outcomes of the 310,000 individual matches, which provides a fuller picture of our model's financial performance, as well as highlights the effects of outliers on our data.

The observed outcomes from the 310,000 simulated matches are summarized in Table 13 and are visualized in Figure 8.

Table 13. Summary Data from 310,000 Simulated Matches.

Variable	GGR	GGR Margin
μ (mean)	$1,531.57	9.09%
σ (standard deviation)	$1,350.55	6.44%
τ (precision)	5.48×10^{-7}	241.41
Min	-$6,290.25	-14.48%
Lower Quartile	$625.07	4.57%
Median	$1,345.29	9.09%
Upper Quartile	$2,247.01	13.78%
Max	$18,184.22	28.87%
IQR	$1,621.94	9.21%

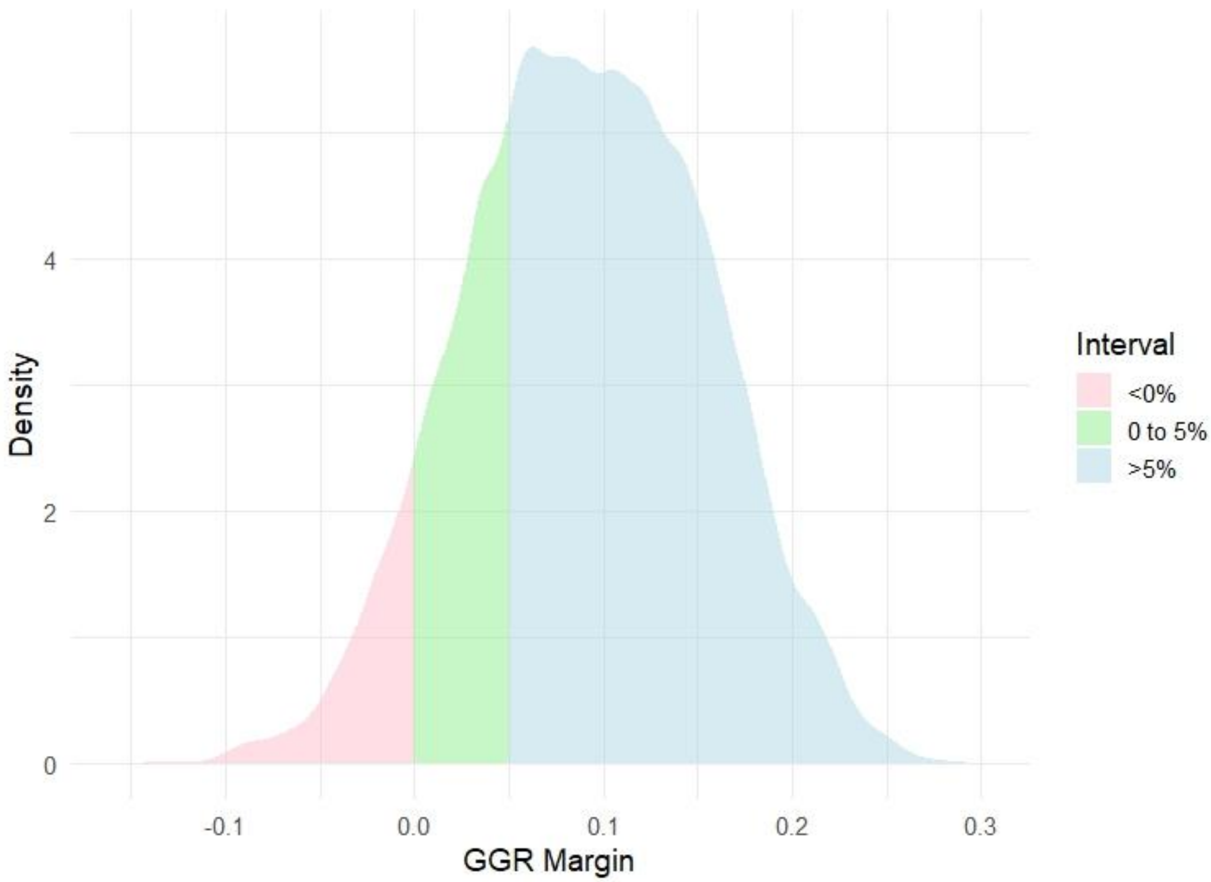

Figure 8. Density intervals for the GGR Margin from the data.

We use the gathered data (from Table 13) to form a new likelihood distribution. While 310,000 simulated matches is an impressive amount of data, these are still only 310,000 theoretical draws for our model. We are still uncertain about the actual financial performance of our model, and thus reform our data into the following likelihood distribution, $p(x|\theta)$, described by our summary (from Table 13):

$$p(x|\theta) \sim \mathcal{C}(0.0909, 0.0461)$$

The real location parameter (x) for this likelihood distribution is the median of our data (0.0909) while the scale factor (γ) is half the interquartile range (0.0461).

Using these revised parameters, we can compare our (revised) likelihood distribution with our prior distribution, as shown in Figure 9.

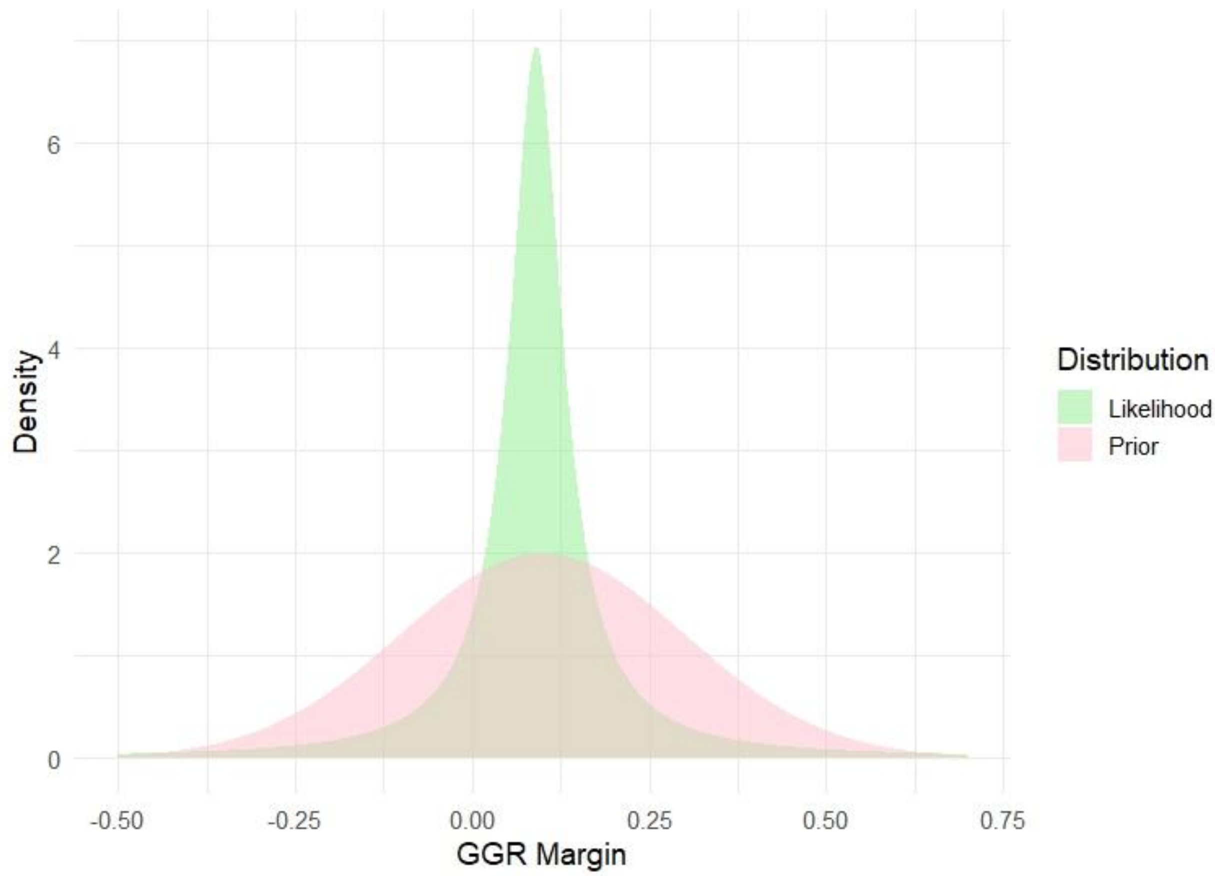

Figure 9. Comparison of the prior and (revised) likelihood distribution for GGR margin.

We move next to form our posterior distribution for the model, derived from both the prior and (revised) likelihood distributions.

C. Posterior Distribution

Our posterior distribution is a function of both our prior and likelihood distributions. Our prior distribution takes the form $p(\theta) \sim \mathcal{N}(0.1000, 0.2000)$ and our likelihood distribution takes the form $p(x|\theta) \sim \mathcal{C}(0.0909, 0.0461)$. We expand the notation to:

$$p(\theta|x) \propto p(x|\theta) \cdot p(\theta), \text{ where}$$
$$p(x|\theta) = \frac{1}{\pi\gamma} \cdot \frac{1}{1 + \left(\frac{x-\theta}{\gamma}\right)^2}, \text{ and}$$
$$p(\theta) = \frac{1}{\sqrt{2\pi\sigma^2}^{e^{-\frac{(\theta-\mu)^2}{2\sigma^2}}}} \quad (25)$$

To solve for the posterior, we simplify to:

$$p(\theta|x) \propto \frac{e^{-\frac{(\theta-\mu)^2}{2\sigma^2}}}{1 + \left(\frac{x-\theta}{\gamma}\right)^2} \quad (26)$$

Equation 26 is intractable, as the solution has no closed standard form. To draw inferences about the potential financial performance of our model, we employ `R` and the `metropolis` function to apply the Metropolis-Hastings algorithm. This algorithm will allow us to draw samples from our open-form posterior distribution by using a Markov chain Monte Carlo (MCMC) process. As shown in the accompanying `R` code, we set the parameters of our prior and likelihood distributions as previously described and sample 5,000 iterations from out posterior distribution [41]. Of note, also, is our choice of starting position from which to draw samples. We set the MCMC process to begin at -0.5 GGR margin to probe for negative outliers first. We chose 5,000 iterations for the algorithm (and not more), to illustrate the effect of potential outliers. As the number of iterations increases, the effect of outliers becomes less and less vivid. By choosing 5,000 iterations, we choose a reasonable number of draws to preserve the desired illustration of the outlier effect.

A visualization of this posterior distribution sampling, along with relevant sampling diagnostics, is shown in Figure 10 and a comparison of the prior, likelihood, and sample draws from the posterior distribution is shown in Figure 11.

The `R` code for this posterior sampling and the accompanying diagnostics is available on the project website [41].

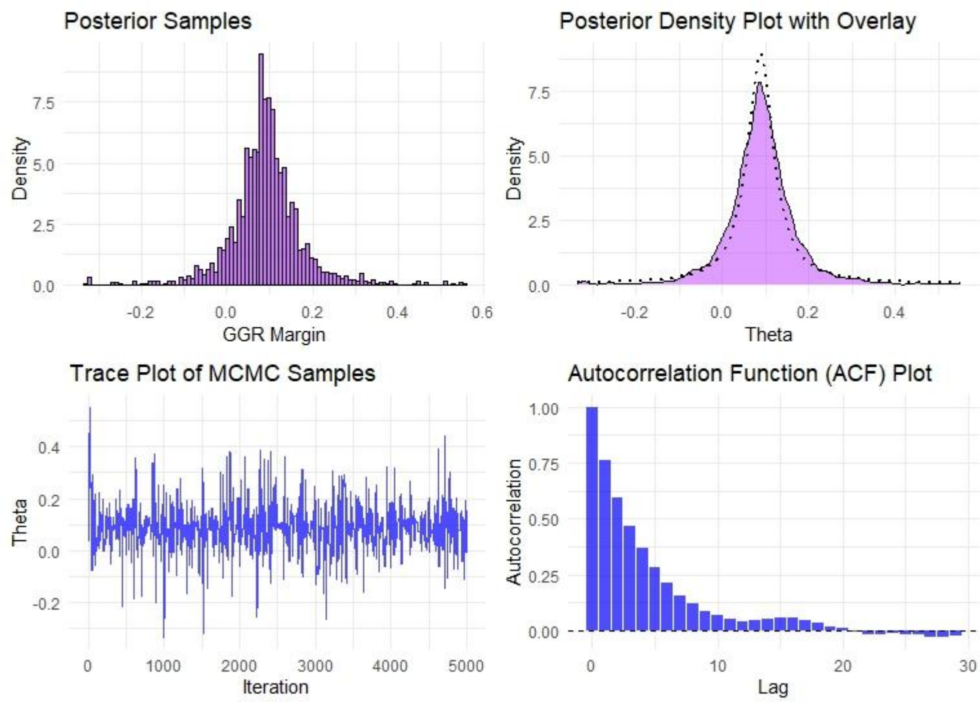

Figure 10. Diagnostic charts from the posterior distribution samples. **Upper left:** visualization of the posterior distribution samples. **Upper right:** visualization of the posterior distribution density and the outlined plot of a Cauchy distribution, $\mathcal{C}(0.0890, 0.0354)$, which loosely approximates the posterior. **Lower left:** Trace plot of the MCMC sampling showing a good "random walk" pattern. **Lower right:** autocorrelation plot showing quick convergence on 0 ACF.

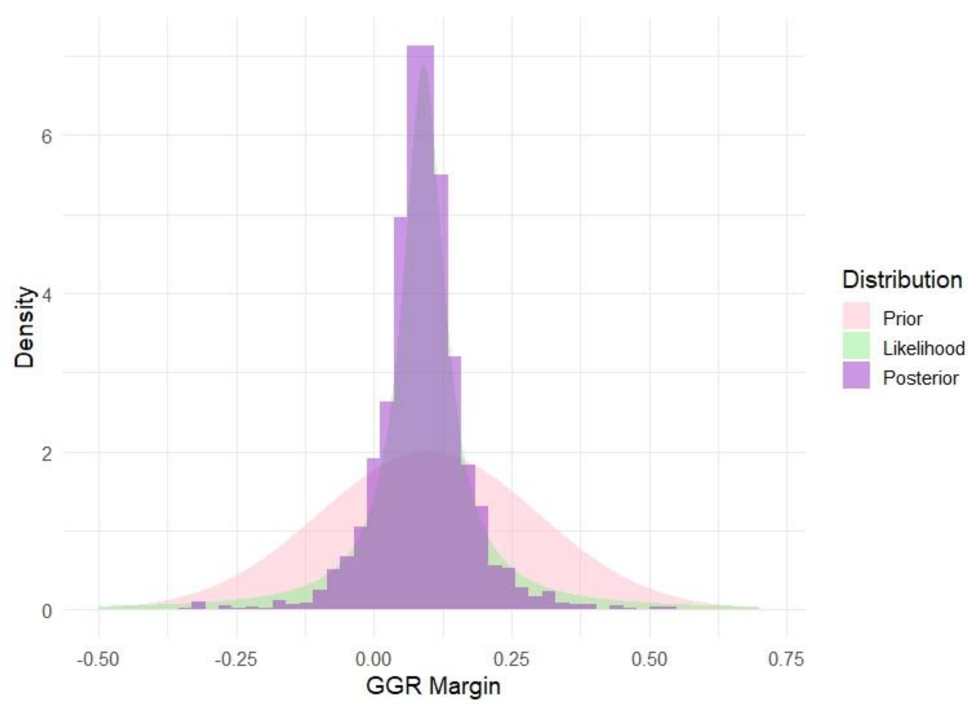

Figure 11. Comparison of the prior and (revised) likelihood distributions alongside the simulated 5,000 sample draws from the posterior distribution.

D. Assessment of GGR Margin

With our posterior distribution calculated, and our 5,000 samples taken from it, we turn now to infer the probability that our model will produce the desired level of GGR margin (that is, >5%). First, we review the summary of the posterior samples, as shown in Table 15.

Table 15. Summary of Posterior Samples.

Variable	Value
μ (mean)	0.0898
σ (standard deviation)	0.0819
τ (precision)	148.9889
Min	-0.3338
Lower Quartile	0.0511
Median	0.0872
Upper Quartile	0.1284
Max	0.5528
IQR	0.0773

We see that, as we expected, there exist large outliers in the posterior distribution. As shown in Table 17, we find minimum results of -33.38% and maximum results of +55.28% GGR margin. These are the largest losses and largest gains, respectively, experienced by the sportsbook from the posterior sampling.

Next, we assess both the GGR margin density from the 5,000 sample draws at the specific intervals shown on Table 16, and the 90 and 95% credible intervals (CI) from those same draws, shown on Table 17.

Table 16. Posterior Probability Density Intervals for GGR Margin.

GGR Interval	Probability Density
< 0%	0.0918
0 to 5%	0.1512
> 5%	0.7570

Table 17. Posterior Probability Credible Intervals for GGR Margin.

CI %	Credible Interval
90%	[-0.0331 , 0.2121]
95%	[-0.0718, 0.2646]

As shown in Table 16, fully 90.82% of the posterior density is profitable (that is, with GGR margin >0%), with 75.7% of the posterior density above our desired level of profit (that is, GGR margin >5%). These density intervals drawn from the posterior density are visualized in Figure 12.

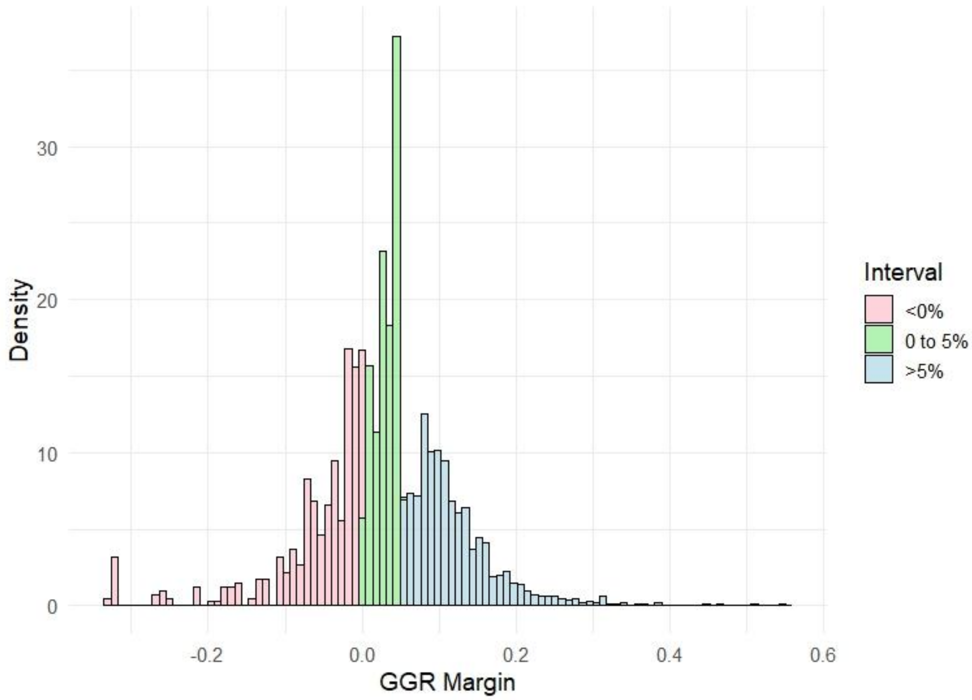

Figure 12. A visualization of the posterior samples at specific intervals. Note the large outliers on the negative side of the x-axis (GGR margin), and the more frequent, but longer "tail" out outliers on the positive side.

Looking at Table 17, we find that the 90% credible interval falls between -3.31 and +21.21% GGR margin and the 95% credible interval falls between -7.18 and +26.46% GGR margin. We thus expect that our model will produce a GGR margin between -7.18 and +26.46% with 0.9500 probability.

Of note, again, are the large outliers. We further calculate that 0.6% of outcomes fall below the -20% GGR margin threshold with correspondingly 6.34% of outcomes above the +20% GGR margin threshold. These outliers are clearly visible in Figure 12.

XI. CONCLUSIONS & FURTHER RESEARCH

This study presents a foundational sportsbook model for competitive trading card games (TCGs), demonstrating its feasibility and profitability through Monte Carlo simulations and Bayesian inference. By integrating Glicko-2 skill ratings, a deck strength assessment model, and structured odds pricing, we developed a system that consistently produces a gross gaming revenue (GGR) margin exceeding the 5% profitability threshold. Our analysis indicates a 75.7% probability of achieving this level of profit, with a 95% credible interval revealing the worst loss on this interval of just 7.18%. The results show some financial risks—namely the 0.6% of outcomes below -20% GGR margin with the single largest sampled loss of -33.38%— but these are mitigated by the model's "risk management" strategy. The implementation of such a system should bear these downside risks in mind and prepare for them accordingly. Overall, results are skewed heavily in the positive space—with 6.34% of outcomes above +20% GGR margin with the single largest sampled gain of +55.28%—reinforcing the model's potential viability for sportsbook applications. We conclude that the model is valid, given the model's assumptions.

While these results provide strong theoretical support for a TCG sportsbook, several challenges remain. Regulatory feasibility, real-world data integration, market demand, and risk management complexities are all critical factors that must be addressed before practical implementation. The lack of existing TCG betting markets means there is little precedent for regulatory approval, and legal barriers in key jurisdictions could pose challenges. Additionally, real-world competitor and deck performance data would improve model accuracy, and bettor behavior analysis is needed to assess potential customer engagement and underlying demand for such a betting product.

Future research should focus on five key areas:

i. Empirical Validation

Use real-world TCG tournament data to test and refine the skill rating, deck strength, and odds pricing models.

ii. Regulatory Assessment

Investigate the legal landscape for TCG betting across major jurisdictions to assess legal feasibility.

iii. Risk Management Enhancement

Explore more sophisticated liability-limiting mechanisms, including dynamic odds adjustments and hedging systems. It is our opinion that *outliers should be adequately planned for*. This means especially *negative outliers*, as these can have severe financial consequences. Future work should continue with a conservative approach, assuming the worst outcomes with high probability (that is, higher than predicted by a Gaussian distribution), to protect against financial disaster for the sportsbook.

iv. Market Analysis & Bettor Behavior

Conduct analyses to better understand bettor behavior, pricing sensitivity, and engagement patterns.

v. Operational and Technological Considerations

Design a framework for data aggregation, real-time odds generation, and model integration with existing sportsbook operational systems.

A potential roadmap toward implementation would involve iterating the model presented in this paper with real-world data, developing risk management and customer engagement strategies, and creating a fully functional sportsbook platform. Continued refinement through testing, regulatory discussions, and market feedback would ensure the model remains viable. A potential "road map" to full implementation might look like that shown in Figure 13.

In conclusion, this paper provides a structured and data-driven approach to exploring TCG betting as an emerging market within the broader "esports" betting landscape. While additional research and validation are necessary, the findings suggest that a well-designed TCG sportsbook could present a viable and lucrative opportunity for operators willing to navigate the associated complexities.

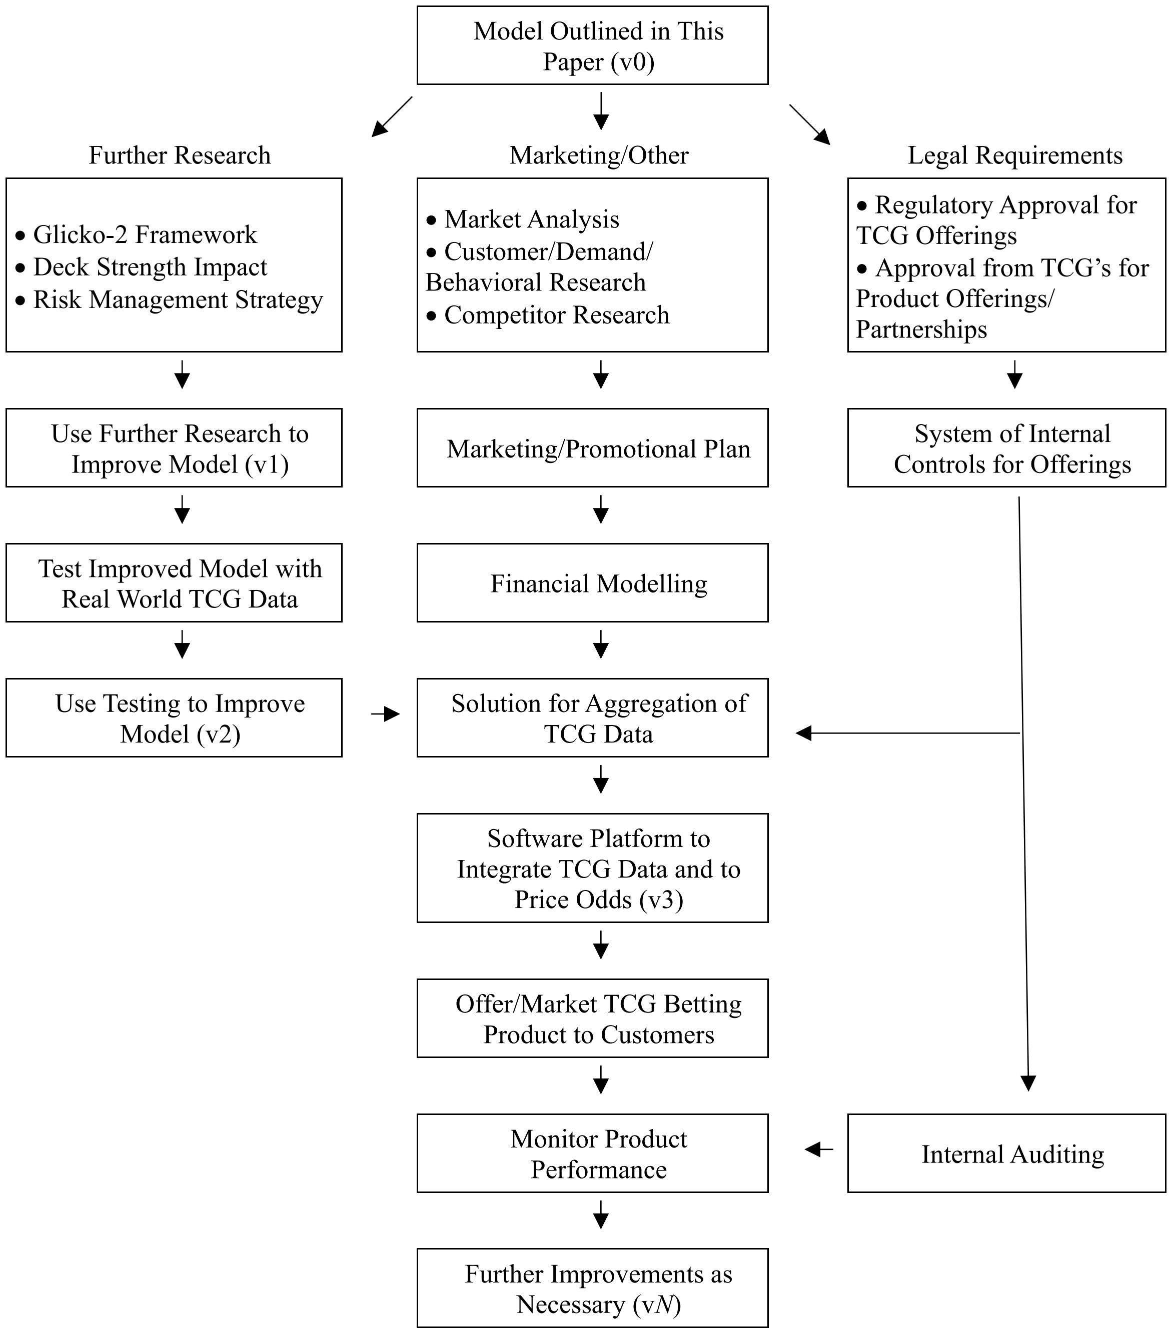

Figure 13. Potential "road map" for future implementation of this TCG sportsbook model. The "road map" sees three new iterated versions (v1 through v3) of the model before launch to the public, with continued improved versions (v*N*) thereafter.

REFERENCES

[1] DraftKings (n.d.). *E-Sports Odds.* Last accessed 25 February 2025. https://sportsbook.draftkings.com/sports/esports

[2] Wizards of the Coast (2023). *Pro Tour Aetherdrift Fact Sheet for Competitors* Magic.gg. Published 19 October 2023. Last accessed 25 February 2025. https://magic.gg/events/pro-tour-aetherdrift-fact-sheet-for-competitors

[3] Pokémon (2024). *2025 Pokémon World Champsionships.* Published 1 November 2024. Last accessed 25 February 2025. https://www.pokemon.com/us/play-pokemon/pokemon-events/championship-series/2025/world-championships

[4] Beaver, N.A. (2025). *Project 1A: An Outline of a Basic TCG Sportsbook Model [Portfolio page].* NicholasABeaver.com. Published 7 March 2025. https://www.nicholasabeaver.com/portfolio/project-1a-an-outline-of-a-basic-tcg-sportsbook-model/

a. *Project 1A: Tournament Season Match Data [.csv file].*
b. *Project 1A: R Code for Glicko-2 Skill Ratings [.txt file].*
c. *Project 1A: Tournament Season Final Ratings [.csv file].*
d. *Project 1A: R Code for Logistic Curve for Rating Deck Strength [.txt file].*
e. *Project 1A: Glicko-2 Ratings for 32 World Tournament Competitors [.csv file].*
f. *Project 1A: Deck Strength Gamma Vector for World Tournament [.csv file].*
g. *Project 1A: R Code for World Tournament Win Probability Matrix [.txt file].*
h. *Project 1A: World Tournament Win Probability Matrix [.csv file].*
i. *Project 1A: TCG Sportsbook World Tournament Simulator Workbook [.xlsx file].*
j. *Project 1A: 10,000 TCG World Tournament Simulations [.csv file].*
k. *Project 1A: World Tournament Simulation GGR and GGR Margin [.csv file].*
l. *Project 1A: R Code for Posterior Sampling and Diagnostics [.txt file].*

[5] Minka, T., Cleven, R., and Zaykov, Y. (2018). *TrueSkill 2: An improved Bayesian skill rating system.* Microsoft. Published March 2018. MSR-TR-2018-8. https://www.microsoft.com/en-us/research/publication/trueskill-2-improved-bayesian-skill-rating-system/

[6] Glickman, M.E. (2022). *Example of the Glicko-2 system.* Boston University. 22 March 2022. http://www.glicko.net/glicko/glicko2.pdf

[7] Petrella, S. (2021). *How Do Sports Betting Odds Work? American Odds, Explained.* Action Network. Published 30 September 2021. Updated 12 February 2025. https://www.actionnetwork.com/education/american-odds

[8] AGA (2019). *Responsible Gaming Code of Conduct.* American Gaming Association. Published 3 January 2019. https://www.americangaming.org/member-code-of-conduct/

[9] Stephenson, A. and Sonas, J. (2020). *PlayerRatings: Dynamic Updating Methods for Player Ratings Estimation* [R package]. Published 1 March 2020. DOI: 10.32614/CRAN.package.PlayerRatings