Quantifying Skill on Opinion Trading Platforms

Amitabha Bagchi
Department of Computer Science and Engineering
Indian Institute of Technology Delhi
bagchi@cse.iitd.ac.in

Ankan Pal
Probo Media Technologies Private Limited
ankan@probo.in

Suyash Kumar
Probo Media Technologies Private Limited
suyash@probo.in

March 24, 2025

Abstract

Opinion trading platforms provide a gamified information market in which users can trade on the outcomes of ongoing events in a range of domains including sports, politics and economics among others. The dramatic increase in popularity of such platforms, and the fact that several such platforms provide real-money gaming options, has led to some discussion on whether opinion trading is a form of gambling. Since jurisprudence on the question of legality of real-money games rests on the question of whether the game is predominantly skill-based or chance based, we demonstrate in this paper that opinion trading is a skill-based game. Through theoretical analysis we establish that skill predominates over chance in opinion trading. We further validate this argument through empirical analysis on real data drawn from the Probo opinion trading platform. Further, we present statistical tests and analytics on Probo data to show that opinion trading shares three other characteristics that games of skill have: Consistency of performance, a skill gradient, i.e., some players are better than others, and a lerning effect, i.e., that players get better over time. The code and data sets used in this work have been made freely available.

Keywords: Information Markets; Opinion Trading; Stochastic Modeling; Mixed Games; Skill in Games; Skill Tests; Skill Scores

1 Introduction

Hahn and Tetlock define an *information market* as a "market for contracts that yield payments based on the outcome of an uncertain future event" [12]. In such a market, unlike in commodity or equity markets, create assets which are claims that will pay off depending on the state of the world at a later point in time. As such information market contracts can be seen as a future markets contract with the asset being traded being non-material. In order to understand how such markets work, let us assume that an England versus Australia cricket match is scheduled for a particular date. An event contract may be created for the outcome of the match. We assume that traders are allowed to buy shares in the two possible outcomes. Eventually when the match ends and the outcome is known the money paid to buy the losing outcome is distributed among the traders who bought the winning outcome. Such information markets are also known as *prediction markets* or *opinion trading markets*. We will use these terms interchangeably.

Electronic information markets have a long and varied history. Beginning with political predictions and sports engagement [12], they have have been used widely within companies as tools for decision making [6]. Their wide adoption led to their being suggested as a framework for public decision making by prominent economists [12, 1]. In recent times prediction markets have proliferated around the world and drawn significant attention from traders engaging with a variety of topics. The recent move by Kalshi to list its event contracts on brokerages is a pivotal moment that is likely to put event contracts on par with equity and commodities [7]. One of the oldest and largest opinion trading platforms in India is offered by Probo Media Technologies Private Limited (Probo) which began operations in 2020 and has a user base of 47 million users as of March 2025. We refer the reader to Appendix A for more details on the Probo platform. The research presented in this paper was funded by Probo and done in close collaboration between IIT Delhi and Probo.

Since the advent of electronic information markets coincided with the advent of online gambling, there has been continued regulatory suspicion of information markets, despite the intervention of reputed economists in their favor [1]. The key question posed by regulators and courts in the context of opinion trading markets is this: Is opinion trading a game of skill or a game of chance? Several decades of jurisprudence and research has address the question of whether a particular real money based activity is a game of skill or a game of chance, a knotty question that arises from the fact that even games that are universally recognized to be skill-based have at least a small component of chance. For example, in Chess it is known that White has a small advantage over Black [25]. In settings such as board games and card games there is a much greater explicit use of chance—the roll of dice or drawing from shuffled decks—and so several efforts have been made over the years to quantify the relative impacts of skill and chance in various settings. In Section 2 we present a detailed discussion of the legal landscape of the skill versus chance question and discuss the different approaches to skill quantification in settings like poker, fantasy sports, board games and mutual fund trading. Based on the approaches established in the literature we identify the following four characteristics of opinion trading that, if established, can support the notion that opinion is a game of skill.

S1. (*Predominance*) A skilled performs better on Probo than a player making random moves.

S2. (*Consistency*) A skilled player performs well consistently over time.

S3. (*Learning*) Active players acquire skill over time in the early stages of their enagagement with the game.

S4. (*Gradient*) The population of players has individuals of different skill levels.

Establishing predominance (S1). In order to address the question of predominance, we take two approaches. In Section 4 we provide rigorous mathematical analysis to demonstrate that players making random moves do not succeed at opinion trading. Since the key domain skill in opinion trading is a good estimate of the answer to the question being traded, we model skill in terms of knowledge of the answer and show that a skilled player is likey to have a high return on investment when pitted against random players. On the empirical side we introduce a novel notion: a skill "dilution" test (Section 5.3). Here we take real event from the Probo platform and with some probability "dilute" the skill of the players by changing the final outcome of the event with some probability. The intuition is that in a game of pure chance a random flip of the eventual outcome should not affect the the players chances of success. We show that in opinion trading such dilution affects the success of players.

Empirical tests for consistency (S2) and learning (S3). In Section 5.4 we show empirically on a real data set drawn from Probo that opinion traders display consistency by measuring the month-to-month correlation of two success metrics: return on investment (ROI) and *win rate* defined as the fraction of trading events in which their return is more than their investment. We establish that there is a learning curve in opinion trading in Section 5.5 by computing correlations between the two success metrics and the number of games played by real players on Probo.

Establishing a skill gradient S4. Following the practice prevalent in skill-based games like Chess, we devise a scoring system we call OpTraS(*Op*inion *Tra*ading *S*core) which incorporates the key skill aspects of opinion trading and the two success metrics ROI and win rate. We then show that the success dynamics of populations of real players on Probo with different OpTraSscores diverge. The scoring system and analysis are presented in Section 6.

Paper organization. We present the legal landscape of the Skill versus Chance debate and discuss the past work in skill quantification in Section 2. In Section 3 we present a mathematical model for an opinion trading platfom. This allows us to present our theoretical analysis of skill predominance in a rigorous manner in Section 4. The empirical analysis for predominance (**S1**), consistency (**S2**) and learning (**S3**) are presented in Section 5. The OpTraSscoring system and the analysis of a gradient (**S4**) in skill level is presented in Section 6. Finally, we discuss the implications of our findings in Section 7.

2 Skill versus Chance

In this section we lay out the context for our study. In Section 2.1 we summarise the legal landscape of real-money gaming and try to locate information markets in this landscape. In Section 2.2 we discuss some of the past work in studying and quantifying skill in real-money gaming.

2.1 Legal Landscape

The distinction between games of skill and games of chance has been a pivotal factor is determining whether a particular real-money wagering activity amounts to gambling or not, c.f., e.g., The Public Gambling Act [11] enacted in 1867 in India. Recognizing that many real-money games contains aspects of both skill and chance, a doctrine of "predominance" was developed by US courts in the context of poker, i.e., several courts held that in the game of poker skill predominates over chance as the factor contributing to a player's success and hence poker cannot be rendered illegal as a gambling activity (c.f., e.g., [32, 5, 33, 26]. In the Indian context the predominance of skill in the card game of rummy was recognized by the Supreme Court in *State of Andhra Pradesh vs K Satyanarayana (1968)* [30], a judgment whose protection has subsequently been extended to online rummy as well [17].

Notably the central legal reasoning based on predominance of skill of *State of Andhra Pradesh vs K Satyanarayana (1968)* [30] was elaborated in the context of horse racing in *Dr KR Lakshmanan vs State of Tamil Nadu (1996)* [31] where the court held that wagering on horse racing requires assessing the form of the horse, jockey, and other variables, distinguishing it from pure games of chance. This argument is similar to the argument made by the New York court of Appeals in *White vs Cuomo (2022)* [21] where it upheld a state law legalizing Daily Fantasy Sports, noting that fantasy game players must draw on their sports knowledge, analyze statistics, and strategically select fantasy team rosters. A similar sentiment has been echoed by Indian courts in the context of cricket-based fantasy sports [3, 13]. Since opinion trading can be seen a more general version of fantasy sport, i.e., it subsumes fantasy sports, the observation, for example, of the court in *Varun Gumber vs Union Territory of Chandigarh (2017)* [13] that such games require "substantial knowledge, attention, judgment, and adroitness" is likely to extend to opinion trading as well.

We refer the reader to Appendix B for a more detailed discussion of the legal landscape.

2.2 Skill Quantification

While there are several methods adopted to quantify skill in the literature, typically these methods fall into three categories: (a) They try to establish some notion of predominance of skill over chance (b) they try to show that skill persists and (c) they try to show that any population has players of different skill level, i.e., that there is a skill gradient. These three categories correspond to **S1** (Predominance), **S2** (Consistency) and **S3** (Gradient) presented in Section 1. We survey some of the literature under each of these themes. We will also briefly comment on the aspect of learning and its relationship to skill, i.e., we try to contextualise **S4** (Learning).

Predominance. A direct approach to testing predominance can be seen in the work of Goodman et. al. [10] who measured the outcomes of several board games with a number of random seeds and demonstrated the varying effect of those seeds on the outcomes. A similar approach leads Banerjee et. al. [2] to conclude that Rummy has a greater predominance of skill than Teen Patti. Misra et. al. [20] study Fantasy Sprots and Mutual Fund trades, proposing a statistical test based on the idea that the outcomes of a game of chance are likely to be concentrated around the mean outcome. Another approach compares actual outcomes to randomized benchmarks: for example, Getty et. al. [9] showed that fantasy sports players consistently beat computer-simulated opponents.

Consistency. Unlike predominance which can be subjective, consistency is easier to measure and so there are numerous studies across domains that argue in favour of skill in a gaming setting by demonstrating consistency of player outcomes. For example Getty et. al. [9] studied Fantasy Sports and Potter et.al. [24] studied poker in this context. It is worth noting though that Potter et. al. [24] were able to establish that a large number of hands needed to be played before skill begins to dominate chance, which can be placed in context against the finding by Meyer et. al. [19] that in shorter sequences of games luck predominates. In a setting very close to our own, Cowgill et. al. [6] examined data from the internal prediction markets of companies like Google and Ford and established a positive, significant correlation between past and future trading profits.

Gradient. A corollary of consistency is a differentiation between players of different skill level. In the context of board games Goodman et. al. [10] found that in some games the presence of chance (e.g. rolls of dice) helped weaker players win occasionally, but stronger players still won more often. Potter et. al. [24] demonstrated that stronger players do better in the long run in poker. Cowgill et. al. [6] were able to separate corporate opinion market traders into better performers who made more profits and those who didn't perform as well.

Learning. In the Psychology literature on skill-based games, it is widely accepted that players learn and get better in such games, c.f., e.g. [27]. This naturally leads to conclusion that such a learning effect should be visible in mixed skill-chance games where skill predominates. Recently Banerjee et. al. [2] have observed this in the context of Rummy. Our work showing such an effect in opinion trading, pushes this line of inquiry further.

3 Modeling Opinion Trading as a Multiplayer Game

Opinion trading platforms enable their user base to trade on questions with Yes/No answers which essentially makes them a multiplayer game platform. In order to analyze the role of skill in this game we first present in overview certain concepts needed to understand the platform dynamics. Subsequently we will define these more formally. The terminology and dynamics presented here closely follow the terminology used in the Probo platform, but all opinion trading platforms work with the same concepts, perhaps naming them differently.

3.1 Overview and terminology

- A *trader*, also referred to in this paper as a "player" or a "user" is a person who uses the platform to trade on Yes/No questions using real money to purchase a stake in either or both answers.
- An *event contract* is the basic object of trade. It contains a Yes/No question along with an expiry time at which its answer is expected to be authoritatively known.

- A *trade* is an action by which a player either offers to purchase (buy) or sell a share in one of the two possible answers. This share is expressed as a discrete number of units that we refer to as the *quantity* of the answer being requested or offered. Each trade also has a *price* at which the offer is being made.
- The *order book* is the list of unresolved trades at a particular time.
- Two trades are said to be *likely to be matched* if one of them offers to buy a particular quantity at a particular price and the other offers to sell that quantity at the same price. At regular interval trades that are likely to be matched are *matched* and removed from the order book, with the appropriate transfer of quantities and money being made.

A system such as the one realized by these concepts is referred to as an *Opinion trading platform* (OTP) or an *Information market*. Some well known opinion trading platforms are Iowa Electronic Market (the earliest one, mainly used for research purposes), Kalshi (CFTC Regulated), Metaculus (forecast and research oriented), and Polymarket (Cryptocurrency-oriented). The originator of this study, Probo, is the largest opinion trading platform in terms of volume, with its user base largely residing in India. Some other India-based platforms focused on opinion trading is sports domains are Sportsbaazi and MPL Opinio.

Notation and Terminology. We denote the set of users by $\mathcal{U}$ and the set of tradable outcomes by $\mathcal{A} = \{\mathsf{Yes}, \mathsf{No}\}$. For an outcome $a \in \mathcal{A}$ we will use the notation $\neg a$ to denote the other outcome, i.e., $\neg\mathsf{Yes} = \mathsf{No}$ and vice versa. Time is considered to be discrete. Each event contract is a tuple $e = \langle e_{\text{id}}, e_{\text{init}}, e_{\text{exp}} \rangle$ where e_{id} is the unique id of the event contract, e_{init} is the time the event contract is initiated and e_{exp} is the time it expires. We denote by $\mathcal{E}$ the set of all event contracts and for any time t we denote by $\mathcal{E}_t$ the set of event contracts that are active at time t, i.e., $\mathcal{E}_t = \{e \in \mathcal{E} : e_{\text{init}} \leq t \leq e_{\text{exp}}\}$. At time t, each user $u \in \mathcal{U}$ has a set $H(u,t)$ of *holdings* which are tuples of the form $h = \langle h_e, h_a, h_q, h_p \rangle$ denoting that u bought a quantity $h_q > 0$ of outcome $h_a \in \mathcal{A}$ for event $h_e \in \mathcal{E}_t$ at price h_p. The presumption is that this was bought at some time $t' \leq t$. We note that each holding corresponds to a single transaction undertaken by the user and for a given event and a given outcome $H(u,t)$ may contain multiple holdings for the same event and same outcome, with possibly differing quantities and differing prices.

3.2 Trading Dynamics in Opinion Trading Platforms

Following the practice of Probo, we will assume that on our OTP the price of the outcomes varies from 0 to 10 in jumps of 0.5, i.e., there are 21 price levels. At each time instant t for each event $e \in \mathcal{E}_t$ both outcomes from $\mathcal{A}$ are associated with a price. We denote the price of outcome a of e at time t by $\mathsf{pr}_{e,t}(a)$. Since the two outcomes of an event are, by definition, mutually exclusive, our OTP implements a coupling of the prices in the sense that at any time t, for any event $e \in \mathcal{E}_t$, $\mathsf{pr}_{e,t}(a) + \mathsf{pr}_{e,t}(\neg a)$ for any $a \in \mathcal{A}$, i.e., the Yes and No prices must add to 10. Clearly this means that the trades in these two outcomes cannot happen independently of each other.

We now discuss the trading dynamics on our OTP. The set of trading actions is denoted $\mathcal{B} = \{\mathsf{Buy}, \mathsf{Sell}\}$. We formally define a trade as follows:

Definition 1 (Trade)**.** *A* trade *is an action initiated by a user to sell or purchase a quantity of an outcome, i.e., a trade is a tuple $s = \langle s_u, s_t, s_e, s_a, s_i, s_q, s_p\rangle$ in which at time s_t, user s_u expresses the intent $s_i \in \mathcal{B}$ for quantity $s_q > 0$ with respect to outcome $s_a \in \mathcal{A}$ for event contract $s_e \in \mathcal{E}_{s_t}$ at price s_p.*

For a holding $h = \langle h_e, h_a, h_q, h_p\rangle \in H(u, t)$ the user u may choose to initiate a trade $\langle u, t, h_e, h_a, \mathsf{Sell}, h_q, h_p\rangle$, i.e., the user may wish to sell that holding. Such a trade is referred to as an exit*. We will refer to trades that are not exits as* ordinary *trades.*

A trade is a statement of intent that needs to be resolved for a transaction to take place. Unresolved trades are stored in an order book as mentioned above. Trades are resolved through a process of matching that we explain below. However there are three other ways in which trades can be removed from the order book: (i) The event expires without the trade being resolved, (ii) the user cancels an unresolved trade, or (iii) the user exercises the "stop loss" action provided by the platform which leads to automatic cancellation of the trade if the price of the outcome being traded moves outside a specified range.

Price coupling of the two trading outcomes is expressed for the purpose of resolving trades as follows: Buying Yes at price p is considered equivalent to selling No at price $10 - p$. We state this formally

Observation 1 (Buy-Sell equivalence)**.** *A trade $s = \langle s_u, s_t, s_e, s_a, \mathsf{Sell}, s_q, s_p\rangle$ is equivalent for matching purposes to the trade $\langle s_u, s_t, s_e, \neg s_a, \mathsf{Buy}, s_q, 10 - s_p\rangle$.*

With this equivalence in hand, we can represent exit trades as Buy trades. Now we are ready to explain how trades are resolved in the system:

Definition 2 (Matching trades)**.** *Two trades r and s with $r_u \neq s_u$ and $r_e = s_e$ are said to* match *if $r_i = \mathsf{Buy}$, $s_i = \mathsf{Buy}$, $r_a = \neg s_a$, $r_p + s_p \geq 10$, contingent on the fact that both r and s are not exit trades.*

Definition 3 (Transaction)**.** *Given two matching trades r and s with $r_u \neq s_u$ and $r_e = s_e$, $r_i = \mathsf{Buy}$, $s_i = \mathsf{Buy}$, $r_a = \neg s_a$, $r_p + s_p = 10$, a* transaction *at time t involves creating a holding $h^r = \langle r_e, r_a, \min\{r_q, s_q\}, r_p$ which is added to $H(r_u, t)$ and a holding $h^s = \langle s_e, s_a, \min\{r_q, s_q\}, s_p\rangle$ which is added to $H(s_u, t)$. Finally the transaction $h = \langle r_u, h^r, s_u, h^s\rangle$ is added to the* transaction register *associated with the event, $\mathcal{T}_{r_e}$. We refer to $\min\{r_q, s_q\}$ as the volume of the transaction h, denoted $\mathsf{vol}(h)$. On adding the transaction to the register the platform deducts $\mathsf{payin}(h, r_u) = r_p \cdot \mathsf{vol}(h)$ from r_u's wallet and $\mathsf{payin}(h, r_u) = s_p \cdot \mathsf{vol}(h)$ from s_u's wallet.*

From Definition 2 we note that there are two kinds of transactions: (a) an exit matched with an ordinary trade and (b) two ordinary trades matched with each other. The key point to note here is that only the second kind of transaction increases the number of entries and the total volume of the transaction register associated with an event e, $\mathsf{Vol}(e) = \sum_{h \in \mathcal{T}_e} \mathsf{vol}(h)$. When an exit is matched with an ordinary trade, an existing entry is just rewritten with the name of the user whose ordinary trade has matched with the exit trade. For example, if the transaction register contained the entry $\langle a, \langle e, \mathsf{Yes}, 12, 7\rangle, b, \langle e, \mathsf{No}, 12, 3\rangle\rangle$ and the user a initiated an exit trade $\langle a, t, e, \mathsf{Yes}, \mathsf{Sell}, 12, 7\rangle$ which matched with user c's trade $\langle c, t, e, \mathsf{Yes}, \mathsf{Buy}, 12, 7\rangle$ then the transaction register entry will get rewritten to $\langle a, \langle e, \mathsf{Yes}, 12, 7\rangle, c, \langle e, \mathsf{No}, 12, 3\rangle\rangle$. One holding will be removed from $H(a, t)$ and one added to $H(c, t)$.

The clearing process. At the time event e expires, (e_{exp}), the event *settlement outcome* o_e is revealed. At e_{exp}, an automated query is triggered to an information source (referred as *source of truth/settlement*) with e_{id} as an input. Corresponding to the rules defined for the event contract e, the information source sends back the settlement outcome $o_e \in \{\mathsf{Yes}, \mathsf{No}, \mathsf{Null\ and\ Void}\}$ as a tuple $\langle e_{id}, o_e \rangle$. If o_e is Yes, the platform takes every entry $h = \langle u_1, \langle e, \mathsf{Yes}, q, p \rangle, u_2, \langle e, \mathsf{No}, q, 10-p \rangle \rangle \in \mathcal{T}_e$ and transfers ₹ $\mathsf{payout}(h, u_1) = 10 \cdot q$ to u_1's wallet. Since u_2 is the loser in this case $\mathsf{payout}(h, u_2) = 0$. Clearance is performed symmetrically in case the settlement outcome is revealed to be No. In case of a $\mathsf{Null\ and\ Void}$ outcome the orders are canceled and the money of the users are reverted back to their respective wallet. For a particular event e the last entry in $\mathcal{T}_e$ corresponds to the clearance. This state of event e is referred to as *settled* and there is no further updation of the transaction register $\mathcal{T}_e$.

Definition 4 (Return on investment)**.** *Denote by $\hat{\mathcal{E}}_t$ all the events that have expired by time t. For a given $u \in \mathcal{U}$ and $e \in \hat{\mathcal{E}}_t$, we define*

$$\mathsf{return}(u, e) = \sum_{h \in \mathcal{T}_e} \mathsf{payout}(h, u),$$

and

$$\mathsf{inv}(u, e) = \sum_{h \in \mathcal{T}_e} \mathsf{payin}(h, u).$$

Then the user u's return on investment *at time t is given by*

$$\mathsf{ROI}(u, t) = \frac{\sum_{e \in \hat{\mathcal{E}}_t} \mathsf{return}(u, e)}{\sum_{e \in \hat{\mathcal{E}}_t} \mathsf{inv}(u, e)}.$$

Definition 5 (Win Rate)**.** *If we denote by $\hat{\mathcal{E}}_t^{(u)}$ all the expired events until time t that the user u has traded in, then we can introduce the notion of **win** for a user u by counting the number of events e in which the $\mathsf{ROI}(u, t) \geq 1$. We can compute the win rate $\mathsf{WR}(u, t)$ as,*

$$\mathsf{WR}(u, t) = \frac{count(\mathsf{ROI}(u, t) \geq 1)}{|\hat{\mathcal{E}}_t^{(u)}|}.$$

4 Mathematical skill analysis

In order to demonstrate that opinion trading on Probo requires skill we model a random trader. We will study the ROI associated with a single event contract. Given the complexity of the trading environment on the platform, we make the following assumptions

- Without loss of generality we will assume that the event outcome is No.
- There are $k + 1$ players.
- The event contract runs for n time steps before expiring.

- A *random player* is one who posts an ordinary Buy trade at each time step where the outcome is selected uniformly at random from $\mathcal{A}$ and the price is selected unformly at random from the $\ell = 19$ outcomes $\{0.5, 1, \ldots, 9.5\}$ independently of the outcome selected.
- The quantity associated with each trade is 1 unit.
- If the trade entered at a time step is not matched it is immediately cancelled.
- There are no exit trades.

4.1 Random player in a random environment

Here we assume that all $k+1$ players are random. The ROI of player 1 can be shown to be below 1 with high probability:

Theorem 1. *Suppose that we have $k+1$ random traders as defined above for some $k > 0$ on the Probo platform. For a single event contract that runs for n time steps, in the limited trading scenario described above, if $\rho \in (0,1)$ is the platform fee proportion deducted from the winnings of the winning traders at the end of the event, there is a constant $c > 0$ depending only on ℓ and k such that at expiry time t*

$$\mathbb{P}\left\{\mathsf{ROI}(\textit{Player 1}, t) > 1\right\} \leq e^{-c\rho^2 n}.$$

Proof. Like all the players, this player enters n trades into the system. For $1 \leq i \leq n$ we denote the price of the ith trade by X_i. To model the outcomes, we define an indicator random variable A_i which is 1 if the outcome being bid for is the winning outcome No and 0 otherwise. Further we denote by E_i the indicator of the event that one of the other k players has entered a trade that matches player 1's trade. So, player 1's investment is

$$Y = \sum_{i=1}^{n} X_i \cdot E_i,$$

and return is

$$Z = 10 \cdot \sum_{i=1}^{n} A_i E_i.$$

Since there are 2ℓ possible trades, $\mathbb{P}\left\{E_i = 1\right\} = 1 - \left(1 - \frac{1}{2\ell}\right)^k$ which we denote by $p(\ell, k)$. Now, given that $\mathbb{E}\left[X_i\right] = 5$ for each i, we can say that

$$\mathbb{E}\left[Y\right] = 5np(\ell, k),$$

and similarly, observing that $\mathbb{E}\left[A_i\right] = 1/2$ for each i,

$$\mathbb{E}\left[Z\right] = 10 \cdot \frac{n}{2} p(\ell, k) = 5np(\ell, k).$$

So we observe that in expectation the random players expected return is the same as the expected investment. Once the platform fee is factored in, the expected return is actually below the expected investment, i.e., the player loses money even though all the other players are also playing randomly without any information. We now study the ROI, i.e., Z/Y.

Using a Chernoff bound for the sum of independent indicators (c.f., e.g., [4]) we get, for $\varepsilon \in (0,1)$

$$\mathbb{P}\left\{|Z - 5np(\ell,k)| > \varepsilon 5np(\ell,k)\right\} \leq 2\exp-\left\{\frac{5\varepsilon^2 np(\ell,k)}{3}\right\}. \tag{1}$$

Since the summands of Y are bounded random variables that take values in the range $[0.5, 9.5]$, Y is a sub-Gaussian random variable with variance bound $\nu = (9.5-0.5)^2 n/4 = 81n/4$, and so we can use Hoeffding's inequality [4] to obtain

$$\mathbb{P}\left\{|Y - 5np(\ell,k)| > \varepsilon 5np(\ell,k)\right\} \leq \exp-\left\{\frac{50\varepsilon^2 np(\ell,k)}{81}\right\}. \tag{2}$$

From (1) and (2), setting $\epsilon = \rho/2$, the result follows by observing that even if the investment is $(1-\rho/2)5np(\ell,k)$ and the return is $(1+\rho/2)5np(\ell,k)$, the deduction of the platform fee, reduces the winnings and therefore the ROI to 1 or lower. □

Theorem 1 directly implies that the expected win rate of a random player is likely to be very low, with the actual fraction being controlled by the length (in terms of steps) of the event contracts entered into, and the platform fee. We state this as a corollary

Corollary 1. *Suppose we have a random player u, as defined above, who participates in a set of contracts that each run for at least n time steps, the last of which ends at time t then*

$$\mathbb{E}\left[\mathsf{WR}(u,t)\right] \leq e^{-c\rho^2 n},$$

where $c > 0$ is a constant that depends on the number of price levels ℓ and the number of players in the event contracts that u participates in.

Theorem 1 demonstrates that a random player stands to end up with an ROI less than 1 with very high probability. In fact the probability of the ROI being favourable to the random player decreases exponentially with the number of trades initiated by the player. However, since all the players are equivalent, one might expect that at least one of them ends up with sizeable winnings. However, this too can be shown to be bounded. We state this fact as a theorem.

Theorem 2. *Suppose that we have $k+1$ random traders as defined above for some $k > 0$ on the Probo platform. If we say that $p(\ell,k) = 1 - \left(1 - \frac{1}{2\ell}\right)^k$ where ℓ is the number of price levels, then, for a single event contract that runs for n time steps, we have that for each $1 \leq j \leq k+1$,*

$$\mathbb{E}\left[\mathsf{return}(\textit{Player } j, t)\right] = 5np(\ell,k)$$

and

$$\mathbb{E}\left[\max_{j=1}^{k+1} |\mathsf{return}(\textit{Player } j, t) - 5np(\ell,k)|\right] \leq 9\sqrt{\frac{n\log(k+1)}{2}}.$$

Proof. We use the fact that for each $1 \leq j \leq k+1$, $\mathsf{return}(\text{Player } j, t)$ is sub-Gaussian with variance bound $\nu = 81/4$ along with the general bound shown on page 31 of [4]. $\square$

Hence we see that return of even the most successful of the random players grows slower than logarithmically in the number of players and as the square root of the number of trades, which is a very slow growth on both accounts.

4.2 Skilled player in a random environment

Here we assume that Players 2 to $k+1$ are random players as defined above. Player 1 is a "skilled" player in the sense that they know that the event outcome is going to be No. We also equip Player 1 with a kind of foresight: at every time step Player 1 can see all the trades entered by Players 2 to $k+1$ and is allowed to initiate the trade that maximizes their ROI.

Theorem 3. *Suppose that we have one skilled trader, Player 1, and k random traders as defined above for some $k > 0$ on the Probo platform. For a single event contract that runs for n time steps, in the limited trading scenario described above, if $\rho \in (0,1)$ is the platform fee proportion deducted from the winnings of the winning traders at the end of the event then at expiry time t*

$$\mathbb{E}\left[\mathsf{return}(\textit{Player 1}, t) > 1\right] \geq (2.1 - \rho)\mathbb{E}\left[\mathsf{inv}(\textit{Player 1}, t)\right].$$

Proof. For $2 \leq i \leq k+1$ and $1 \leq t \leq n$, let $X_{i,t}$ be the price of the trade entered by Player i and $A_{i,t} = 1$ if the outcome is Yes and 0 otherwise. Let $Y_t = \max_{i=2}^{k+1} X_{i,t} \cdot A_{i,t}$ and let $I_t = \max_{i=2}^{k+1} \cdot A_{i,t}$. Player 1 will be interested in matching a trade only if the outcome bid for is Yes (since Player 1 knows that the event outcome is No). At every time step t, Player 1 will enter a trade for outcome No with price $Z_t = 10 - Y_t$.

Player 1's investment is $S = \sum_{i=1}^{n} I_t Z_t$ and return is $R = 10 \sum_{i=1}^{n} I_t$. Since I_t is 0 only if all the k random players want to buy the outcome No, the expectation of the return is

$$\mathbb{E}\left[R\right] = 10 \sum_{i=1}^{n} I_t = 10n\left(1 - \frac{1}{2^k}\right).$$

To simplify the analysis of the investment we define the following collection of random variables: For each $2 \leq i \leq k+1$ and $1 \leq t \leq n$ let

$$H_{i,t} = \begin{cases} 0 & \text{w. prob. } 1/2 \\ r \in \{0.5, 1, \ldots, 9.5\} & \text{w. prob. } 1/2\ell \end{cases}$$

With this definition we note that $I_t Z_t$ has the same distribution as $J_t = \max_{i=2}^{k+1} H_{i,t}$ for each t. For simplicity we bound $\mathbb{E}\left[2J_t\right]$ by observing that for $1 \leq r \leq 19$,

$$\mathbb{P}\left\{2J_t \geq r\right\} = 1 - \left\{\frac{1}{2} + \left(\frac{r-1}{2\ell}\right)^k\right\} = \frac{1}{2} - \left(\frac{r-1}{2\ell}\right)^k.$$

Therefore

$$\mathbb{E}\left[J_t\right] = \sum_{r \geq 1} \mathbb{P}\left\{J_t \geq r\right\} = \frac{1}{2}\left\{\frac{\ell}{2} - \sum_{r=1}^{\ell}\left(\frac{r-1}{2\ell}\right)^k\right\} \leq \frac{\ell}{4} = \frac{19}{4}.$$

And so $\mathbb{E}\left[S\right] \leq 19n/4$ and the result follows. □

Discussion. Theorem 1, Corollary 1, Theorem 2 and Theorem 3 together provide mathematical evidence to support the hypothesis **S1**, i.e., that skill predominates over chance on an OTP.

4.3 Extending the mathematical model to more realistic scenarios

The mathematical model studied above has two key limitations. Firstly, it avoids dealing with exit trades and secondly it works with a simplistic and unrealistic model of skill: it assumes the skilled trader already knows the event outcome. Other limitations in this model are that (i) quantities traded are kept fixed to 1, (ii) every trader is expected to trade at every time step. We discuss some extensions of the simple model analyzed above but leave the actual analysis for future work.

Extension 1: Evolving knowledge scenario. In this extension we have 1 skilled trader and k random traders but the skilled trader doesn't have absolute knowledge of the event outcome. We model the skilled trader's evolving knowledge by assuming that at time t the skilled trader guesses the correct outcome with probability p_t. p_1 is assumed to be $1/2$ and $p_{t+1} > p_t$. In this case we try and prove theorem similar to Theorem 3 to estimate how well the skilled player can do in such a setting.

Extension 2: Hierarchical knowledge scenario. Here we have k players and each of them is similar to the skilled player of Extension 1. The hierarchy is established by ordering their probability of guessing the correct outcome. In such a scenario we will estimate the ROI of players at different levels of the hierarchy.

Extension 3. Exiting with skill. In the evolving knowledge scenario a trader can realise at some point that an earlier buy was wrongly entered and may want to exit that holding. So we can enrich the hierarchical knowledge scenario by allowing exit trades and then again estimate the ROI of players at different levels of the hierarchy.

Insights. We note that the analysis contained in Theorem 1, Eq. (2), and Theorem 3 contains the basic ideas that make us feel that these extensions will also yield similar insights, i.e., that random players are unlikely to do well, even if all other players are playing randomly, and that a skilled player is likely to earn a high ROI in the system if the other players have lower skill levels.

5 Skill tests for Opinion Trading Platforms

In this section we present a set of skill tests on real data drawn from the Probo platform in order to argue that skill plays a significant role in opinion trading.

5.1 Overview of statistical skill tests

The Skill Dilution Test. Skill predominance **S1** is studied through this test. Here we take a set of real event contracts and modify their outcomes randomly. We choose a randomisation parameter $\alpha \in [0, 1]$ and for each event e with outcome o_e we flip the outcome to $\neg o_e$ with probability α. The payout at clearance is modified accordingly. We then study the ROI and win rate of the users in the period under study for different values of α. The idea here is that trader skill is expressed by their ability to correctly estimate the outcome of an even. By randomising the outcome we are effectively "diluting" any skill that is inherent in the trader actions. A degradation of trader performance implies that there was some skill inherent in those actions and randomising event outcomes did in fact dilute that skill.

Persistence Test The time consistency in the performance of skilled users **S2** is studied through the *persistence test*. For each month of the calendar year 2024, we consider a group of users whose monthly trading volume exceeds 20. We compute the ROI and WR of these users for each month and study the *correlation* between the performance parameters for each pair of consecutive months.

Learning Test This test considers the learning aspect **S3** which establishes that skill level increases with active usage of the platform. We consider a group of users who have traded in at least 360 events in the year 2024. We study the correlation of the mean and median WR and ROI of the users with the number of events played (which we also refer to as *event rank*).

5.2 Data sets

We work with 3 data sets drawn from the Probo platform, we call them Probo-1, Probo-2 and Probo-3.

Probo-1. In this data set we consider the performance (ROI and WR) of a set of $125, 986$ *unique* users whose number of orders are ≥ 50 across $256, 453$ events in a 90 day time window between 30 Nov 2024 and 28 Feb 2025.

Probo-2. In this data set we consider the performance (ROI and WR) of a set of $418, 960$ *unique* users whose number of trades are ≥ 50 across $1, 136, 182$ events in the whole year 2024. As we considers the performance parameters month-wise, we provide the month-wise split-up of the number of users in Table 2.

Probo-3. In this data set we consider the performance (ROI and WR) of a set of $37, 242$ *unique* users who registered on or before 1 Jan 2024 and who have played at least 360 events in the 360 day time window between 28 Jan 2024, 21 Jan 2025.

Access. The data sets are freely available for download and the code can be accessed through this GitHub Repository.

Table 1: Data Set Details

Data Set	Time Period	No. of Users	No. of Events	Restrictions
Probo-1	30 Nov 2024 to 28 Feb 2025	125,986	256,453	Number of Orders ≥ 50
Probo-2	1 Jan 2024 to 31 Dec 2024	$418,960$	$1,136,182$	Number of Orders ≥ 20
Probo-3	28 Jan 2024 to 21 Jan 2025	37,242	720	Users who registered on or before 1 Jan 2024

Table 2: Probo-2 Data Set Monthly Split up

	Jan-Feb	Feb-Mar	Mar-Apr	Apr-May	May-Jun	Jun-Jul	Jul-Aug	Aug-Sep	Sep-Oct	Oct-Nov	Nov-Dec
No. of Users	71,557	65,132	82,045	162,091	182,913	113,947	79,653	57,635	47,828	48,318	52,335
No. of Events	218,487	212,350	178,576	164,925	183,282	193,337	192,518	189,086	182,506	179,324	180,629

5.3 The Skill Dilution Test for Opinion Trading Platforms

Assuming that the beginning of the data collection for Probo-1 was at time 0 and the end time was T we denote the set of completed events by $\hat{\mathcal{E}}_T$, and the set of outcomes of those events by $\mathcal{O} = \{o_e : e \in \hat{\mathcal{E}}_T\}$. We discard those events for which $o_e =$ Null and Void. We denote the set of users in data set Probo-1 by $\mathcal{U}$. We name the success metric vectors of all the users in $\mathcal{U}$ by $\mathsf{ROI} = \{\mathsf{ROI}(u, T) : u \in \mathcal{U}\}$ and $\mathsf{WR} = \{\mathsf{WR}(u, T) : u \in \mathcal{U}\}$ and denote the mean and median of these collections by **Mean** (ROI) , **Median** (ROI) , **Mean** (WR) and **Median** (WR).

Now we pick $\alpha \in [0, 1]$ and generate $\mathcal{O}^{(\alpha)} = \{o_e^{(\alpha)} : e \in \hat{\mathcal{E}}\}$ where

$$o_e^{(\alpha)} = \begin{cases} \neg o_e & \text{with prob. } \alpha \\ o_e & \text{with prob. } 1 - \alpha \end{cases}$$

For each choice of α we update the transaction registers for each $e \in \mathcal{E}_T$ to generate new success metric vectors ROI^{α} and WR^{α}.

Hypothesis testing to check dissimilarity of key statistical parameters under skill dilution. We performed Welch's parametric t-test [34] and Mann-Whitney's [18] non-parametric u-test and to test the basic null hypotheses that the statistics **Mean** (ROI) , **Median** (ROI) , **Mean** (WR) and **Median** (WR) are lesser than or equal to their corresponding statistics under α dilution, i.e. **Mean** $\left(\mathsf{ROI}^{(\alpha)}\right)$, **Median** $\left(\mathsf{ROI}^{(\alpha)}\right)$, **Mean** $\left(\mathsf{WR}^{(\alpha)}\right)$ and **Median** $\left(\mathsf{WR}^{(\alpha)}\right)$ respectively for different values of α. We note that the null hypothesis was rejected in each case. The order of the p-values obtained during the hypothesis testing were $< 10^{-100}$ and hence considered insignificant, therefore we do not report the results in detail. k

Visualizing the effect of skill dilution. We ran α dilution for $\alpha \in [0, 1]$ in steps of 0.1. In Fig. 1 we plotted the mean of the user win rate versus α. Error bars are plotted in the figure but are very close to the curve, showing that the

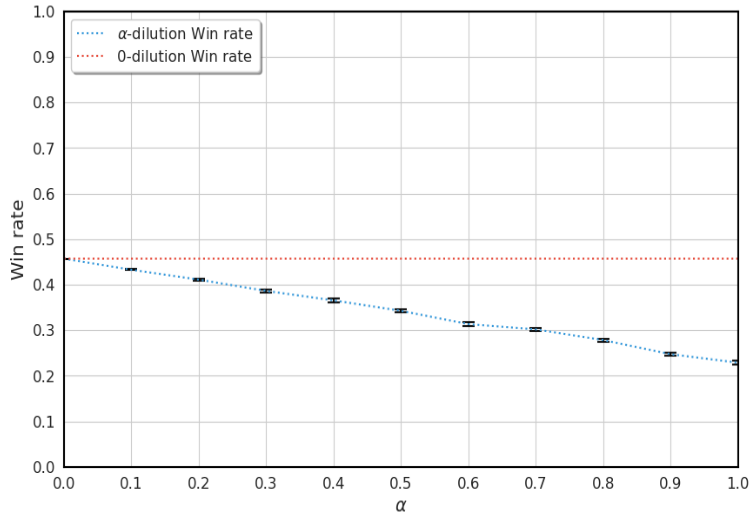

Figure 1: Skill Dilution Test: Mean Win rate Vs α

variance of the mean win rate is low. We can see that as α increases, both statistics show a decrease. This drop in win rate demonstrates that dilution reduces the win rate of users and so provides evidence of the predominance of skill in opinion trading. The entire win rate distributions under dilution are visualized in Fig. 2. We can see here that the entire population of users suffers a decrease in win rates under dilution.

5.4 Persistence Test for Opinion Trading Platforms

In order to test for persistence of success from month to month, we created dataset Probo-2 by aggregating monthwise data from Probo for the calendar year 2024, only choosing users who have at least 20 orders. We compares the ROI and WR for pairs of consecutive months.

Hypothesis testing to establish correlation from month to month. We took the null hypothesis that the Spearman Correlation [14] between the ROIs and Win Rates of users for consecutive months of 2024 is non-positive. The hypothesis was rejected for each pair of months. The order of the p-values obtained during the hypothesis testing were $< 10^{-100}$ so we omit them here. This establishes that there is a month-to-month positive correlation for both the success metrics.

Consistency heatmap. In Fig. 3 we present a heatmap of pairwise correlation coefficients for pairs of months, not just consecutive pairs. We notice that the correlation between $\mathsf{WR}_{Jan\ 2024}$ and $\mathsf{WR}_{Dec\ 2024}$ is 0.59. This implies that the consistency of performance has persisted for the whole period. The correlation varies in $[0.52, 0.65]$ indicative of the high positive correlation between the win rates of users in two different months and subsequent months for the year 2024. We observe a drop in the correlation coefficients in the month of March, April, and May. On Probo, one of the largest category is *Cricket* in terms of total traded volume. In the month of March, April, and May one of biggest cricketing event **IPL** (Indian Premiere League) takes place and hence new users are added who are beginners and their

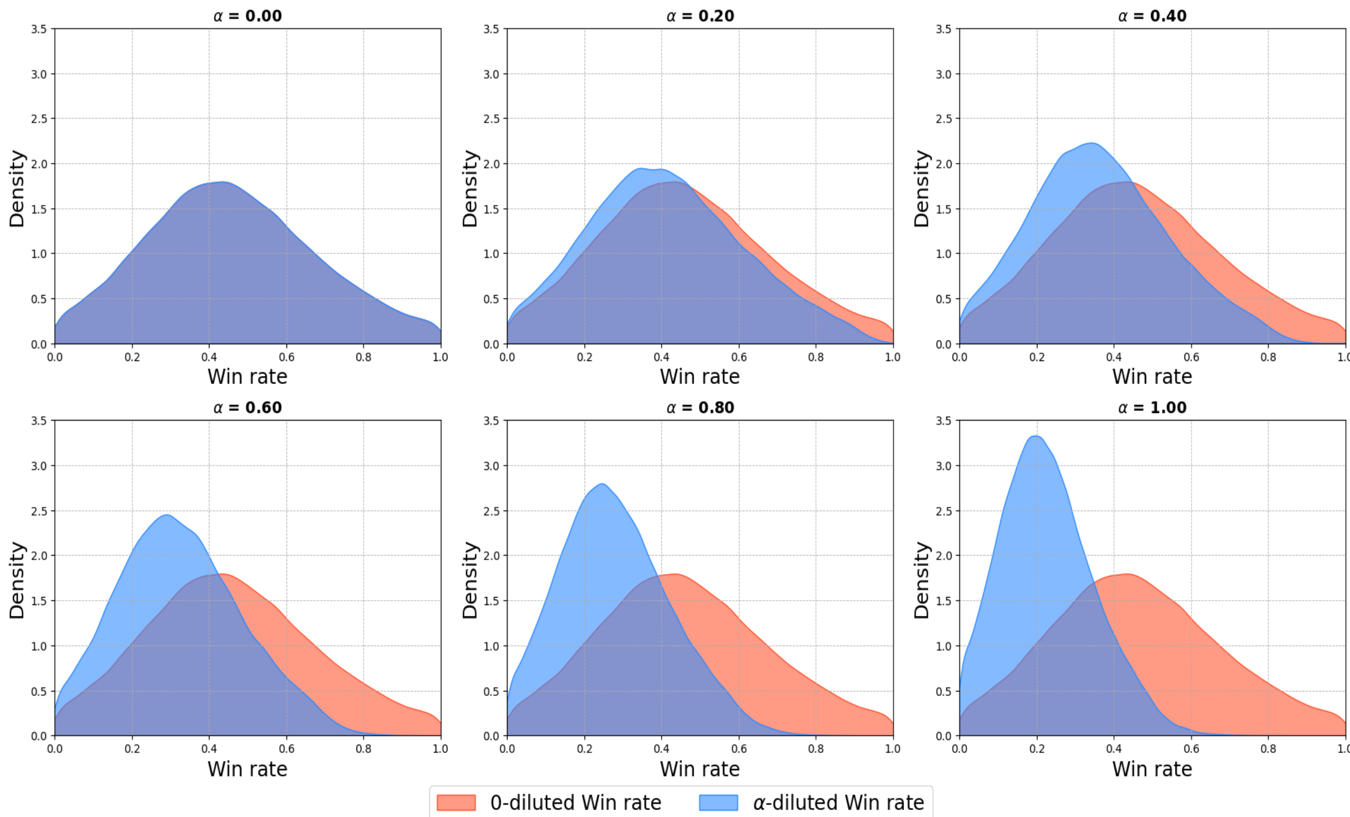

Figure 2: Comparison of Undiluted Win Rates vs α-Diluted Win Rates

performance is usually *inconsistent*. Moreover, due to addition of new users on the platform, the considered sample size for these months is significantly greater than the other months.

5.5 Learning Test for Opinion Trading Platforms

We investigated the aspect of learning in opinion trading. To do so we created a data set Probo-3 of around 37K users, $\mathcal{U}$, who had traded for a volume of at least 20 units in 720 events. For each $1 \leq i \leq 720$ and $u \in \mathcal{U}$ we denote by $e_{u,i}$ the ith event in the sequence of 720 events played by user u. Note that the ith event for two different users may be different. For the purpose of this experiment we defined a version of ROI here based on the event rank i:

$$\mathsf{ROI}(u,i) = \frac{\sum_{j=1}^{i} \mathsf{return}(u, e_{u,j})}{\sum_{j=1}^{i} \mathsf{inv}(u, e_{u,j})}.$$

Similarly we have for each $1 \leq i \leq 720$

$$\mathsf{WR}(u,i) = \frac{count(\mathsf{ROI}(u,i) \geq 1)}{i}.$$

Hypothesis testing to establish a learning effect. For each $1 \leq i \leq 360$ we computed the mean and median of the collections $\{\mathsf{ROI}(u,i) : u \in \mathcal{U}\}$ and $\{\mathsf{WR}(u,i) : u \in \mathcal{U}\}$. For each of the four rank-indexed collections of statistics

Figure 3: Heatmap of Correlation of Win Rates between Subsequent Months

our null hypothesis was that the statistic is not positively correlated with rank using the Spearman Correlation as a measure. The hypothesis was rejected in each case. Since the order of the p-values obtained during the hypothesis testing was $\sim 10^{-100}$, we do not report them here.

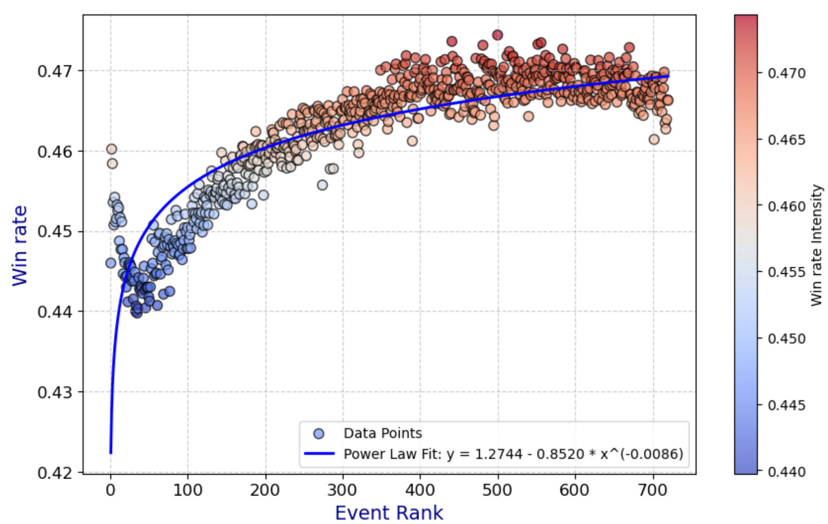

Figure 4: Median Win rate versus Event Rank

Analyzing the learning curve. In Figure 4 we plot **Median** $(\{\mathsf{WR}(u, i) : u \in \mathcal{U}\})$ versus event rank i for $1 \leq i \leq 720$. The curve follows a power law which is similar to the curves seen in the literature associated with the study skill-based games (c.f., e.g., [27, 28, 23].) We studied the rate at which the median win rate changes with event rank for different groups of players, grouping based on their success level by win rate: top $1\%, 10\%, 25\%, 50\%, 75\%$ and top 100%, i.e., all the players. We fitted a power law $y = u - ax^{-c}$ on the curves for each of these groups, where u can be interpreted as the asymptotic (final achievable) performance, a can be interpreted as the initial learning gain/performance level, and c as the learning rate. The slopes acx^{-c-1} are plotted in Figure 5. The top performers

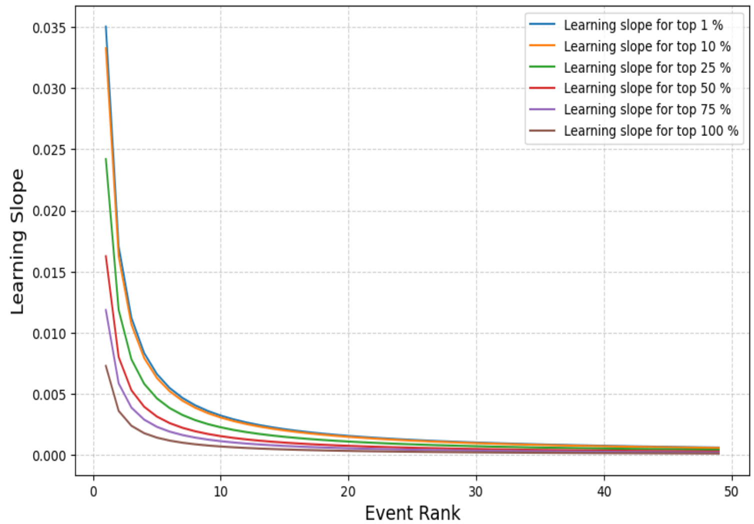

Figure 5: Slope Comparison for Different Performance Levels

learn faster although for all groups there is a plateauing of the learning effect, a phenomenon also observed in the study of skill-based games (c.f., e.g. [27].) In Table 3 we present the values of the slopes observed for the different groups at event ranks 10, 20 and 50. To contextualize these numbers we note that Steyvers and Benjamin [28] studied the slopes of three brain training games on the brain training site Lumosity. The slopes observed across three different games for the youngest group of players who engaged with the platform for the longest time were 0.109 (for *Lost in Migration*), 0.165 (for *Ebb and Flow*) and 0.426 (for *Memory Match*). These slopes were observed early in play but after most early dropouts had left. Comparing these rates for our top 1% of playes we see that at event rank 10 the slope is 0.322 and at event rank 20 the slope is 0.157. Since the games on Lumosity are explicitly designed to develop skill, the comparability of the slopes of the growth in win rate in opinion trading is a powerful argument in support of the claim **S3**, i.e., that there is a learning effect in opinion trading, and, therefore that opinion trading is a skill-based game.

6 The OpTraS scoring system and skill gradient

The OpTraS scoring system seeks to rate the skill level of players by rewarding them for their success, their engagement and their adroitness. Success is represented by two metrics: net money earned, i.e., their returns, and their ROI, which

Table 3: Power law fit parameters for win rate curves at different performance levels. Slopes are scaled by 100.

Percentile ranks	initial learning a	learning rate c	final performance u	slope at rank 10	slope at rank 20	slope at rank 50
Top 1%	1.027	0.034	1.663	0.322	0.157	0.061
Top 10%	1.004	0.033	1.526	0.307	0.15	0.058
Top 25%	0.95	0.025	1.444	0.224	0.11	0.043
Top 50%	0.903	0.017	1.373	0.147	0.072	0.028
Top 75%	0.878	0.013	1.324	0.11	0.055	0.021
All	**0.852**	**0.008**	**1.274**	**0.067**	**0.033**	**0.013**

is the return scaled by the investment. Their engagment is represented by the number of event contracts in which they have participated. And their adroitness is quantified by the number of exit trades they enter into, i.e., the number of times they divest themselves of their holdings before the expiry of an event.

The OpTraS system has been implemented on Probo. All users who have made at least 20 trades receive an OpTraS score. We describe the system below, also giving the values of the systems parameters as optimised for Probo.

Notation and key components of OpTraS. Given a user u at time t we denote the expired events in which u has been active in reverse chronological order of their expiry time by $\{e_{u,t,i}\}_{i=0}^{n_u^t-1}$ where n_u^t is the number of events of $\hat{\mathcal{E}}_t$ in which u has been active. Since we will be applying an exponential decay to the success metrics, we define the rate of decay $\lambda_{u,t}$. On Probo we set this value to $\frac{\ln 2}{0.3 n_{u,t}}$ since this decay rate gives 30% of the recent orders 50% weightage. The rationale behind choosing this rate is that through data analysis on Probo we found that more than 80% of the consistent users who have been trading at least once a day for the whole year of 2023, had 30% new trading actions in a span of 15 days.

Given any metric α defined over any $\mathcal{U}$, we will use $\mathsf{rank}(\alpha, u)$ to denote the rank of $u \in \mathcal{U}$ under that metric with the lowest rank being assumed to be 0. We will use $\mathsf{rankmax}(\alpha)$ to denote $\max_{u \in \mathcal{U}} \mathsf{rank}(\alpha, \mathcal{U})$. Note that it appears that $\mathsf{rankmax}(\alpha)$ should be $\mathcal{U}$ but in practice the number of active users keeps changing so $\mathsf{rankmax}$ can be a dynamic value for any metric. We also denote the percentile rank of $u \in \mathcal{U}$ w.r.t metric α by $\mathsf{prank}(\alpha, u) = 1 - \mathsf{rank}(\alpha, u)/\mathsf{rankmax}(\alpha)$.

We now define four metrics that will be combined to form the OpTraS score.

1. *Performance score.*
$$\pi_u^t = \sum_{i=0}^{n_u^t-1} e^{-i\lambda_{u,t}} \mathsf{return}(u, e_{u,t,i}).$$

2. *Strategy score.* We consider the sequence of eventwise ROI's for u: $\left\{\frac{\mathsf{return}(u,e_{u,t,i})}{\mathsf{inv}(u,e_{u,t,i})} : 0 \le i < n_{u,t}\right\}$. Let $A_{u,t}$ be the weighted average of this sequence weighted by $\{e^{-i\lambda_{u,t}} : 0 \le i < n_{u,t}\}$ and let $M_{u,t}$ be the median of the sequence weighted by the same weights. Then we define the strategy score to be
$$\varrho_u^t = \frac{A_{u,t} + M_{u,t}}{2}.$$

3. *Activity score.* Noting that n^t defines a metric over $\mathcal{U}$, we define the activity score as
$$\theta_u^t = \frac{n_u^t}{1 + e^{-\nu\{\mathsf{prank}(n^t,u) - 0.5\}}},$$
where the sigmoid parameter ν is set to 13 on Probo. This setting ensures that the top 20% of players (in terms of events entered into) lose only 2% of the scored accrued to them by n^t. This helps them overcome transient effects like having a bad week in which they lose a lot of money.

4. *Foresight score.* This score is denoted ϕ_u^t and is simply equal to the fraction of all trades made by u till time t that are exit trades.

The OpTraS score.

Definition 6 (The OpTraS score)**.** *Let us assume $f(x) = 799x + 200$. Then the OpTraS score for user $u \in \mathcal{U}$ at time t is given by*

$$\Lambda_u^t = a_1 f(\mathsf{prank}(\pi^t, u)) + a_1 f(\mathsf{prank}(\varrho^t, u)) + a_1 f(\mathsf{prank}(\theta^t, u)) + a_1 f(\mathsf{prank}(\phi^t, u)),$$

where a_1, a_2, a_3, a_4 are preset constants.

On Probo we have used $a_1 = a_2 = 0.4$ and $a_3 = a_4 = 0.1$.

Using OpTraS to establish skill gradient on Probo. The OpTraS scoring system was implemented on the Probo platform in June 2024. In order to show that there are groups with different skill levels we extracted the OpTraS scores of two groups of users: those who engage consistently with the platform and those who don't.

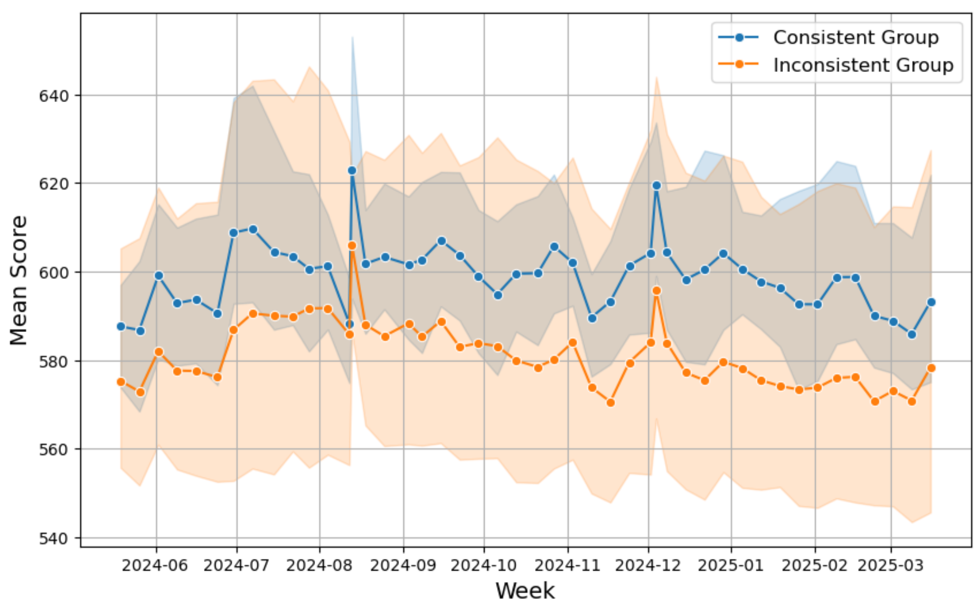

Figure 6: Skill Score Trends for Multiple Category of Users on Probo

Anecdotal evidence within Probo suggested that consistent engagers have a higher skill level. We identified all users who made at least 20 trades a week as *consistent* players and the rest were deemed *inconsistent.* We plotted the OpTraS score of these two groups for a period of 9 month (Figure 6). As we see the consistent group outperformed the inconsistent group consistently. This demonstrates **S4**, i.e., there are groups of players with very different skill levels.

7 Discussion

We have endeavored to demonstrate in this paper that opinion trading has a very significant aspect of skill. In our study we have presented both theoretical and empirical arguments to establish the predominance of skill. We have also shown that skill persists and can be measured in player performance. The existence of a distinct learning curve in player performance, similar to that observed in games that are acknowledged to be skill-based, is a further argument in favor of the claim that opinion trading requires skill, as is the existence of diversity in the skill levels of players.

Although our study has focused on data from Probo, the experimental framework can be applied to any opinion trading platform. We have made the code for the OpTraSscoring system available so that other opinion trading platforms can use it. We also feel that if Fantasy sports platforms wish they can incorporate features from OpTraSinto their own scores. Through this paper and by making the code and data for our studies available we wish to initiate a larger and more inclusive discussion on the legality of opinion trading. Following and extending the arguments made by Arrow et. al. [1], it is our feeling that opinion trading can help players develop their critical and analytical faculties and give them a reason to be better informed about the world. In the long term this can help develop a more informed and engaged citizenry. Our hope is that this work helps opinion trading emerge from the shadow of suspicion and claim its rightful place.

References

[1] K. J. Arrow, R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J. O. Ledyard, S. Levmore, R. Litan, P. Milgrom, F. D. Nelson, G. R. Neumann, M. Ottaviani, T. C. Schelling, R. J. Shiller, V. L. Smith, E. Snowberg, C. R. Sunstein, P. C. Tetlock, P. E. Tetlock, H. R. Varian, J. Wolfers, and E. Zitzewitz. The promise of prediction markets. *Science*, 320(5878):877–878, 2008.

[2] T. Banerjee, A. De, S. Maitra, and D. Mukherjee. Skill vs. chance quantification for popular card & board games. *arXiv preprint arXiv:2410.14363*, 2024.

[3] Bombay High Court. Gurdeep Singh Sachar v. Union of India. PIL No. 16, 2019.

[4] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

[5] Colorado Supreme Court. Colorado v. Raley, 2009.

[6] B. Cowgill, J. Wholfers, and E. Zitzewitz. Using prediction markets to track information flows: Evidence from google. Technical report, Harvard University, 2008. Retrieved on 2 March 2025 from https://www.eecs.harvard.edu/cs286r/courses/fall10/papers/GooglePredictionMarketPaper.pdf.

[7] M. J. de la Merced. Kalshi, an online prediction market, will open its bets to brokerages. New York Times, 31 January 2025. retrieved on 2 March 2025 from https://www.nytimes.com/2025/01/31/business/dealbook/kalshi-prediction-market-brokerages.html.

[8] Federation of Indian Fantasy Sports. The Supreme Court once again affirms fantasy sports offered by Dream11 as a game of skill. `https://fifs.in/the-supreme-court-once-again-affirms-fantasy-sports-offered-by-dream11-as-a-game-of-skill/`. Accessed: 2025-03-17.

[9] D. Getty, H. Li, M. Yano, C. Gao, and A. Hosoi. Luck and the law: Quantifying chance in fantasy sports and other contests. *Siam Review*, 60(4):869–887, 2018.

[10] J. Goodman, D. Perez-Liebana, and S. Lucas. Measuring randomness in tabletop games. In *2024 IEEE Conference on Games (CoG)*, pages 1–8. IEEE, 2024.

[11] Government of India. The Public Gambling Act, 1867. Act No. 3 of 1867.

[12] R. Hahn and P. Tetlock. Using information markets to improve public decision making. *Harvard Journal of Law and Public Policy*, 29(1):214–289, 2005.

[13] High Court of Punjab and Haryana. Shri Varun Gumber v. Union Territory of Chandigarh and Others. CWP No. 7559, 2017.

[14] M. Hollander, D. A. Wolfe, and E. Chicken. *Nonparametric statistical methods*. John Wiley & Sons, 2013.

[15] Illinois Supreme Court. Dew-Becker v. Wu. IL 124472, 2020.

[16] Karnataka High Court. All India Gaming Federation v. State of Karnataka. W.P. No. 18703/2021, 2022.

[17] Kerala High Court. Gameskraft Technologies Pvt. Ltd. v. State of Kerala. W.P.(C) Nos. 7785, 7851, 7853, 8440, 2021.

[18] H. B. Mann and D. R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *Annals of Mathematical Statistics*, pages 50–60, 1947.

[19] G. Meyer, M. von Meduna, T. Brosowski, and T. Hayer. Is poker a game of skill or chance? a quasi-experimental study. *Journal of Gambling Studies*, 29:535–550, 2013.

[20] V. Misra, D. Shah, and V. Ranganathan. Is it luck or skill: establishing role of skill in mutual fund management and fantasy sports. *Massachusetts Institute of Technology*, 2020.

[21] New York Court of Appeals. White v. Cuomo. NY Slip Op. 01954, 2022.

[22] New York Supreme Court. People v. DraftKings. 21 Civ. 5739, 2015.

[23] A. Newell and P. S. Rosenbloom. Mechanisms of skill acquisition and the law of practice. In *Cognitive skills and their acquisition*, pages 1–55. Psychology Press, 2013.

[24] R. J. Potter van Loon, M. J. van den Assem, and D. van Dolder. Beyond chance? the persistence of performance in online poker. *PLoS one*, 10(3):e0115479, 2015.

[25] J. Rowson. *Chess for Zebras: Thinking Differently About Black and White*. Gambit Publications, 2005.

[26] South Carolina Supreme Court. Town of Mt. Pleasant v. Chimento. 737 S.E.2d 830, 2012.

[27] T. Stafford and N. Vaci. Maximizing the potential of digital games for understanding skill acquisition. *Current Directions in Psychological Science*, 31(1):49–55, 2022.

[28] M. Steyvers and A. S. Benjamin. The joint contribution of participation and performance to learning functions: Exploring the effects of age in large-scale data sets. *Behavior research methods*, 51:1531–1543, 2019.

[29] Supreme Court of India. RMD Chamarbaugwala v. Union of India. AIR 1957 SC 628, 1957.

[30] Supreme Court of India. State of Andhra Pradesh v. K. Satyanarayana. AIR 1968 SC 825, 1968.

[31] Supreme Court of India. K.R. Lakshmanan v. State of Tamil Nadu. AIR 1996 SC 1153, 1996.

[32] United States District Court for the District of Nevada. Baxter v. United States. 633 F. Supp. 912, 1986.

[33] United States District Court for the Eastern District of New York. United States v. DiCristina. 886 F. Supp. 2d 164, 2012.

[34] B. L. Welch. The generalization of 'Student's' problem when several different population varlances are involved. *Biometrika*, 34(1-2):28–35, 1947.

A The Probo Opinion Trading Platform

We delve into the platform level statistics of Probo to understand the adoption and acceptance in the Indian market. Since the inception of Probo in 2020, the total number of registered users are $47,233,714$ ($47.23\ million$) until March 2025. Out of these registered users $33,666,303$ ($33.67\ million$) users have traded at least once amounting to $1,371,817,524$ ($1.37\ billion$) orders on $2,836,669$ ($2.84\ million$) events spread over 54 categories. Till mid-March 2025, this activity accounted for total of $160,457,994,760$ ($160.46\ billion$) trades amounting to Gross Merchandise Value (GMV) of ₹ $649,840,873,005.56$ (₹ $649.84\ billion$). To understand the growth of the platform we provide the year-on-year growth trajectory here.

Table 4: Platform Level Statistics of Probo Platform

Year	Total Events	Total Users	Total Orders	Total Trades	Total GMV
2020	1,712	815	23,145	16,970	85,307
2021	18,035	1,947,913	14,563,706	59,901,422	296,344,262
2022	402,371	10,125,544	185,618,072	5,608,644,264	28,445,267,445
2023	1,056,680	11,866,372	479,233,904	38,005,313,278	208,627,124,361
2024	1,142,791	17,702,274	559,143,613	66,943,956,642	344,595,178,506
2025 (Jan-Mar)	216,650	5,605,240	133,227,209	49,837,349,648	67,876,873,126

B The Legal Landscape of Real-Money Gaming

B.1 Overview of Key Legal Cases on distinguishing Skill vs Chance in Online Real Money Games

The distinction between games of skill and games of chance has been a pivotal factor worldwide in legal determination whether a gaming activity amounts to gambling. We will discuss some of the notable cases that have explored this distinction in skill-based games such as opinion trading, poker, daily fantasy sports, and rummy etc. These cases underscore the complex legal landscape and the ongoing legal debates surrounding the classification of games as either skill-based or chance-based, significantly impacting their regulation across various jurisdictions.

B.1.1 Daily Fantasy Sports

FanDuel and DraftKings, leading platforms in daily fantasy sports (DFS), have been at the center of legal debates concerning whether their offerings are games of skill or chance. In November 2015, New York Attorney General argued that these fantasy contests were games of chance rather than skill, making them unlawful. In June 2016, New York passed legislation legalizing and regulating DFS, effectively resolving the legal dispute by classifying these contests as games of skill [22]. Similarly, the New York court of Appeals in *White vs Cuomo* [21] upheld a state law legalizing DFS and noted that fantasy game players must draw on their sports knowledge, analyze statistics, and strategically select fantasy team rosters – activities requiring judgment and expertise and thus predominantly a skill-based activity. The Illinois supreme court applied the Dominant Factor Test to concluded that DFS involves a significant degree of skill, influencing the legal status of such games in Illinois [15].

In India, courts have similarly held fantasy sports to be skill-based. In *Varun Gumber vs Union Territory of Chandigarh (2017)* [13], the Punjab and Haryana High Court ruled that the DFS offered by Dream11 (India's largest DFS platform) is a game of skill, not chance, emphasizing that success depends on participant judgment in team selection rather than on any single match outcome. The high court also asserted that fantasy sports games were skill-based, requiring substantial knowledge, attention, judgment, and adroitness [13]. The Bombay High Court echoed this in *Gurdeep Singh Sachar vs Union of India (2019)* [3], finding "no betting or gambling is involved" in Dream11 because results do not depend on a particular team winning but on players' overall performance, which requires skillful selection. The Supreme Court of India has repeatedly affirmed these findings – dismissing challenges and noting that fantasy sports are "games of skill and a legitimate business activity" protected under the Constitution (Federation of Indian Fantasy Sports) [8].

B.1.2 Poker

In the offline era there was a widespread discussion on whether organizing Poker amounted to gambling and *how* the earnings shall be taxed. In *Baxter vs United States (1986)* [32], the court recognized poker as a game where skill predominates over chance, thereby treating Baxter's earnings from Poker as business income. In Colorado vs Raley (2009) [5] the jury acquitted Raley who organized poker tournaments in Colorado, based on evidence that poker is a game of skill. In 2012, multiple courts in the U.S. ruled that poker is predominantly a game of skill rather than chance

[33] [26]. These landmark decision concluded that under U.S. law, poker does not constitute gambling under the Illegal Gambling Business Act. The rulings relied on extensive expert evidence and concluded that Texas Hold'em poker is dominated by skill – citing factors like strategy, bluffing, and statistical odds – and thus did not qualify as "gambling", post a "predominance test or dominant factor test".

B.2 Indian Legal Landscape around Skill vs Chance in Real Money Gaming

The Supreme Court of India addressed the legality of the card game Rummy, ruling that despite elements of chance in the initial card distribution, the game's outcome is predominantly determined by the player's skill in memorizing and strategizing (*State of Andhra Pradesh vs K Satyanarayana (1968)* [30]). Citing *Satyanarayana*, the apex court held that horse racing is a game of skill, as it requires assessing the form of the horse, jockey, and other variables, distinguishing it from pure games of chance in *Dr KR Lakshmanan vs State of Tamil Nadu (1996)* [31]. The *Satyanarayana* ruling set a broad precedent in India: if a game's outcome is dominated by skill, it is protected as a legitimate business activity and not subject to gambling bans. Indian courts have consistently followed this reasoning. Various high courts and district courts in India have applied the "preponderance of skill" test to many real money games of mixed nature, even in their online variants. In 2021, the Kerala High Court struck down a ban on online rummy [17], affirming that online rummy (whether played for stakes or not) remains a game of skill and is protected by the *Satyanarayana* precedent [30]. Thus, rummy is firmly established in Indian jurisprudence as a skill game – the element of random card dealing is considered insufficient to turn it into gambling, since a *skilled player can consistently outperform an unskilled one in the long run*. Outside India, rummy has not been as heavily litigated, mostly because it is generally seen as a casual game.

One of the notable legal cases in this space is the *All India Gaming Federation (AIGF) vs State of Karnataka* [16] which challenges the constitutionality of the Karnataka Amendment Act No. 28 (2021) which criminalizes the playing or facilitating of online games. AIGF won the case and the amendment was declared unconstitutional and therefore removed. One of the key arguments that the case highlighted was that Skill based activities are *res extra commercium* and protected under Article 19(1)(g) of the constitution of India referring to the *RMD Chamarbaugwalla vs The Union Of India (1957)* [29].