International Review of Law & Economics 74 (2023) 106119

Contents lists available at ScienceDirect

International Review of Law & Economics

journal homepage: www.elsevier.com/locate/irle

The dominance of skill in online poker[☆]

Jerome Hergueux[a,b,*,1], Gabriel Smagghue[c,1]

[a] French National Center For Scientific Research (CNRS, BETA lab), Strasbourg, France
[b] ETH Zurich, Center for Law and Economics, Zurich, Switzerland
[c] Bank of France, Paris, France

ARTICLE INFO

JEL classification:
K20
L83
Z2

Keywords:
Online poker
Skill
Learning
Selection
Gambling

ABSTRACT

Does skill dominate luck in online poker? In many countries around the world, the legality of the online poker industry rests on courts' evaluation of the "skill dominance" criterion. Because it is not precisely defined, however, the skill dominance criterion may be misleading when it comes to the legal qualification of online gambling activities. We argue that this concept might be better framed as "do skilled players dominate the game" than as "does skill dominate game outcomes". We introduce a novel, comprehensive dataset on online poker play – where we follow 91,439 players over 40 consecutive months (representing over 85 million hands played) – and develop simple tests to show that (i) skill in the game drives individual results, and (ii) players improve their skills with experience and quit playing the game as a function of starting ability. A lower bound estimate suggests that it takes at least 7 months of full-time training for a novice to acquire the basic skills exhibited by the most experienced players in our data. We conclude that the scholarly debate around this industry may move beyond that of its legality to focus instead on issues of regulation. Beyond the case of online poker, the procedures and tools we develop can be readily transferred to evaluate the skill dominance criterion in other purported games of skill, such as sports betting or stock trading.

1. Introduction

Poker is one of the most popular card games worldwide, and the question of its legal status poses itself in similar terms in many jurisdictions around the globe. Millions of players play the game every day for real money, either physically or online, representing a multi-billion dollar industry with its class of superstar players, specialized magazines, TV channels, websites, forums and dedicated international competitions (e.g., the World Series of Poker). In an attempt to put the game on par with traditional sports, its proponents have argued that poker should be considered a "mind sport" alongside chess, bridge or Go.[2]

Whether or not poker formally qualifies as a (mind) sport, its trajectory as an industry is not unlike that of "e-sports" (Hallmann and Giel, 2018). The advent of the Internet has spurred a genuine boom in both interest (demand side) and practice (offer side). This, in turn, has given rise to a self-sustained class of professional players (McCormack and Griffiths, 2012; Meng-Lewis et al., 2021), suggesting an important role for skill in the game involving a subtle combination of mathematical ability, self-control and strategic thinking. Similar to e-sports, the current practice of (online) poker can therefore be seen as a manifestation of "sportification" (Heere, 2018). The concept of sportification refers to the idea that some activities that may not be considered a "sport" at first glance (e.g., because of a lack of physicality) may still be analyzed and organized as such, notably when they are played competitively and followed by a large audience. In theory, this should imply a scientifically-informed conversation around how to best organize this activity so that it resembles a sport by allowing a fair, pleasurable,

☆ We are grateful to Adda52 for sharing their data and for their interest in better understanding their players. We are indebted to Charles Nesson, Patrick Nally, Mohit Agarwal, Pratap Pandit and Stefan Bechtold for their support. We thank Kuba Straszewski, Emanuel Schmid, Uwe Schmitt, Roskar Rok and Cristian Scurtescu from the ETH Zurich Scientific IT services, as well as Arturo Vivas, Shiyi Li and Tianqi Wang for research assistance. We acknowledge financial support from the Center for Law and Economics at ETH Zurich, Switzerland, the ETH Zurich Career Seed grant, Switzerland, and the University of Strasbourg Attractivity grant, France.

* Corresponding author at: French National Center For Scientific Research (CNRS, BETA lab), Strasbourg, France.
E-mail addresses: jerome.hergueux@gess.ethz.ch (J. Hergueux), gabriel.smagghue@banque-france.fr (G. Smagghue).
[1] The authors contributed equally to this work.

[2] For instance, the recently created International Federation of Match Poker (see https://matchpokerfed.org/) has sought the recognition of the Court of Arbitration for Sport (CAS) and obtained a status of observer within the International Mind Sports Association (IMSA) and the Global Association of International Sports Federations (GAISF).

https://doi.org/10.1016/j.irle.2022.106119
Received 29 June 2022; Received in revised form 8 December 2022; Accepted 9 December 2022
Available online 14 December 2022

and safe environment for individuals to compete and compare their performance over time.

From a legal perspective, however, the debate around the sportification of online poker never made it to the above regulatory issues, as it stumbles upon a more basic issue: that of it is very legality. The legal question of interest here is that of whether online poker can be distinguished from a game of chance – i.e., online gambling. Indeed, in many legal systems, gambling activities are either prohibited altogether, or fall under a highly supervised legal regime (e.g., state ownership or indirect control via a concession).[3] However obvious it may have appeared to the game's practitioners, the question of whether online poker should be legally considered a game of skill thus emerged as a high stake economic one.

In the U.S., a number of courts arrived at the conclusion that because it was "dominated by chance", online poker should be legally equated to online gambling. The notion of "skill dominance" has not been precisely defined by judges, however, so that we can only speculate as to their underlying mental model by looking at their actual decisions. In a detailed account of the legal debate around this issue, Miles et al. (2013) document how courts ruled in a methodological vacuum that may have precisely resulted from the ambiguity of the concept of skill dominance. Neither did judges seek to define skill in the context of the game, nor did they try to come-up with a methodology to assess its role in determining player performance. Among other examples, decisions of a similar nature have been observed in India – an emerging but fast growing market for online poker – where several high courts have rules in the same direction (Sayta, 2012).

The legal consequences were fierce for the American online poker industry as it fell under the 2006 Unlawful Internet Gambling Enforcement Act (UIGEA). In 2011, the three most important online poker platforms (PokerStars, Full Tilt Poker and Absolute Poker) were indicted on charges of bank fraud, money laundering, and illegal gambling. Their domain names were seized, restraining orders were issued against dozens of bank accounts in the U.S., and the Department of Justice sought over $3 billion in penalties and forfeitures (it eventually settled for over $700 million).

In sharp contrast to this legal assessment, the scientific literature which emerged in parallel has largely converged to the conclusion that (online) poker was a game of skill.[4] Interestingly, almost none of these papers operationalized the "skill dominance" criterion put forth by judges in terms of whether skill "dominated" luck in determining overall game outcomes. Instead, researchers set out to test whether differences in skill between players were "significant", i.e., whether said differences were large enough to enable sizeable differences in performance and profits over time. Croson et al. (2008) collected data on high-profile poker and golf tournament rankings over several years. They found that players' yearly rankings are persistent over the years, and that this persistence is as strong in poker as it is in golf – a game thought to be primarily skill-based. Miles et al. (2013) discussed the issue of skill measurement from a legal perspective and relied on online poker data to show that (i) players differ in their profitability and (ii) their returns tend to be correlated over time. Levitt and Miles (2014) showed that players identified *a priori* as being highly skilled (e.g., because they were top money winners in a previous tournament or appeared in one of the published lists of the best contemporary players) achieve an average return on investment of over 30 percent in the World Series of Poker, compared to −15 percent for all the other players. van Loon et al. (2015) relied on online poker data to show that players who rank in the top (respectively, bottom) decile of earnings in the first half of their sample period (covering 12 consecutive months) are substantially more likely to end up in the top (respectively, bottom) performance decile in the second half. More recently, Duersch et al. (2020) have nuanced the above literature by pointing out that while the case for skill in poker appears indisputable, this does not necessarily imply that game outcomes are *predominantly* determined by skill. Based on a comparison with the game of chess, the authors conclude that online poker outcomes are, in fact, predominantly determined by chance.

In this paper, we argue that the skill dominance criterion understood as "does skill dominate game outcomes" can be highly misleading when it comes to making a legal distinction between gambling and skill-based activities. The reason is straightforward. It puts uncertain decision making environments at a disadvantage in the examination process, as the criterion requires them to pass some arbitrary cutoff in terms of how much of individual outcomes can effectively be attributed to differences in skill. Such an approach represents a significant departure from the way skill has been traditionally characterized across industries in the literature. In many industries, skilled decision making is typically made under stringent conditions of uncertainty that make it virtually impossible for any individual player to "control" the majority of market outcomes. This is why "skill dominance" in such environments is usually defined with respect to the question of whether differences in skill can be detected across players that generate consistent profits. In other words, the criterion asks: "do skilled players dominate the game?"

As a field example, consider the global market for CEOs. Most researchers would not contend with the fact that players in this industry are engaged in a highly competitive game of skill (Gabaix and Landier, 2008). Superior leadership abilities can lead to efficiency gains, visionary strategic decisions and improved employee motivation, all of which creates significant value. At the same time, empirical research has shown that corporate performance as a whole (and, hence, CEO compensation) may be largely determined by random factors (Bertrand and Mullainathan, 2001). As another case in point, consider the example of stock trading. The empirical literature in economics has long recognized that the overwhelming majority of individual investors do not show evidence of skill in stock picking (Barber and Odean, 2000, 2013). On average, the performance of those investors is effectively as good as (and sometimes even worse than) random. Does that imply that individual stock trading should be considered a game of chance — akin to online gambling? Such a position would ignore the fact that skill in stock picking does allow a small number of elite investors in the population to consistently beat the market (Akepanidtaworn et al., 2022).

This paper illustrates those ideas by assembling a novel dataset from comprehensive first-hand play records shared by the first, dominant online poker platform in India.[5] One unprecedented feature of our data is that we can track the decisions and outcomes of *all* No Limit Hold'em players playing for real money from the inception of the platform in January 2015, and up to May 2018 — including the decision to stop playing altogether (i.e., exit decisions). All in all, our dataset includes 91,439 players whom we follow during a period of 3 years and 4 months (representing over 85 million hands played).

[3] For a brief history of the legal status of gambling in various legal systems, see Spapens (2014).

[4] In this section, we only review studies which use field data to conduct their analyses. A small literature also exists which has developed simplified poker games to study related questions in an experimental context. In particular, Meyer et al. (2013) and Van Essen and Wooders (2015) track a number of direct skill measures in their experiments which are closely related the ones we use in this paper. Due to significant differences in designs, however, those experimental results remain difficult to compare. Overall, based on a horizon of 60 hands, Meyer et al. (2013) find weak evidence for skill in poker, while based on a horizon of 200 and 720 hands, respectively, Van Essen and Wooders (2015) and DeDonno and Detterman (2008) report strong evidence in support for skill. We remain cautious about interpreting these experimental results in the field context of our paper, where real-play data can be collected to answer our research question of interest.

[5] See https://www.adda52.com/.

Specifically, we use this data to make two distinct contributions to the literature. Our first contribution is to clarify the legal debate around the ambiguous concept of "skill dominance" put forward by courts in order to legally distinguish games of chance from skill-based activities. We do this both conceptually and empirically. At the conceptual level, we distinguish between two possible interpretations of the concept, understood either as (i) "does skill dominate game outcomes", or as (ii) "do skilled players dominate the game". Both interpretations may very well lead to opposite conclusions, as the former has a (questionable) mechanical tendency to treat uncertain fields less favorably. In practice, we argue that the skill dominance criterion might be better framed empirically as "do skilled players dominate the game" than as "does skill dominate game outcomes". Defined in such a way, the concept does not depend on the level of uncertainty in the environment considered and remains consistent with the way skill is usually characterized in other decision making fields.

We derive the empirical consequences of this legal distinction by developing a simple framework in which both interpretations can be tested independently. Specifically, we use a fixed effect regression to show that (i) skill in our population of players "only" explains about 17.2% of the variance in hand-level outcomes, while (ii) our model is able to reject the hypothesis that skill does not drive player performance with as little as 7 hands of play history. Depending on the selected cutoff for skill dominance, our results illustrate that both tests may indeed lead to contradictory legal conclusions.

Second, we contribute to the advancement of our knowledge on skill dominance in online poker by being the first to work with simple but *direct* measures of skill in the game. While the extant literature has built a convincing case for skill in poker based on indirect evidence (such as player-level persistence in tournament rankings), working with such measures allows us to provide direct evidence on how skillful play impacts individual earnings. It further allows us to describe the skill acquisition process in the population of players. This, in turns, provides the analyst with a simple tool that can be used to understand the game's ecosystem, and draw a less abstract picture of how skill relates to individual results. Specifically, we identify three basic skills that matter for player performance: (i) self-control, (ii) an ability to take (calculated) risk, and (iii) good probability calculation and hand selection.[6]

In a final step to our analysis, we leverage the above skill variables to describe the game's learning ecosystem. Learning within the population of players can take one of two forms: (i) within-player learning, where players improve their skill level as they become more experienced with the game, and (ii) between-player learning, where some players realize that they are less able than others and thus selectively drop-out of the platform as a function of ability. Empirically however, any trend in playing style according to experience will reflect the combination of both effects. We therefore propose a simple model aimed at separating them out.

This above test is of interest to the analyst, as is indicative of a well-functioning competitive environment. Intuitively, gifted players should be more likely to pursue training and improve their skills, while less-able ones should be more likely to stop playing altogether. If learning does not occur along one or the other dimension, this might raise a red flag for the regulator. The game might be rigged (e.g., some players are acting on private information), or losing players might get severely addicted to it. We find that learning occurs at a significant rate across both dimensions: players quit playing poker as a function of starting ability, and surviving players learn to improve their skills as they play more hands. A lower bound estimate suggests that it takes at least 7 months of full-time training for a novice to acquire the basic skills exhibited by the most experienced players in our data.

The rest of the paper proceeds as follows. We provide some general background on the game of poker and derive our main skill constructs of interest in Section 2. We present our dataset and variables in Section 3 and develop our identification strategy in Section 4. We report our empirical results in Section 5, where we perform our skill dominance tests, document the impact of skill on player-level performance, and describe the online poker learning ecosystem. We discuss the legal implications of our results in Section 6.

2. Defining skill in online poker

While the extant literature has relied exclusively on indirect measures of skill to perform its tests (by relying on, e.g., player rankings), our own approach requires that we define a number of indicators of skillful play in the game that can be used as proxies for skill in our analyses. Those proxies need not exhaust the notion of skill in online poker, but they have to provide a solid basis for skillful play, at least in theory.

At a conceptual level, the game of poker which we study is relatively simple. Each poker hand starts with the random distribution of two private cards to each player at the table — the "hole cards". Five "community cards" are then dealt for all players to see, in three stages. All players seek the best five card combination of the community cards and their two hole cards. The first stage consists of a series of three community cards ("the flop"), then an additional single card ("the turn"), and a final card ("the river"). Rounds of betting take place before the flop is dealt and after each subsequent deal. At each stage, players have the option to fold their hand (i.e., quit, leaving any money invested so far on the table), check/call (i.e., match whatever investment the preceding player made to stay in the hand), or raise (increase the amount currently at stake). The player who has the best hand and has not folded by the end of all betting rounds wins all of the money bet for the hand, known as the pot.

Popular accounts of the game of poker sometimes describe it as a game where (the detection of) deception and bluffing represents the predominant skill determining long-term performance. While quite widespread, poker experts consider those perceptions as mostly wrong. Especially in online poker where physical cues are altogether absent (but also in face to face situations), textbooks insist that poker is first and foremost an investment game (Sklansky and Malmuth, 1999; Sklansky and Miller, 2006). To become profitable, players are taught to think in terms of the expected value of their hand, and to invest accordingly. The game has a significant uncertainty component, since the hole cards are dealt privately to each player. As a result, players form an expectation regarding the value of their hand based on its intrinsic strength, but also based on their estimate of the strength of the cards of the other players, which they reveal progressively through their own investments. Assessing this expected value is, therefore, a quite complex decision making problem, which has recently attracted significant attention in the computer science community (Bowling et al., 2015; Moravčík et al., 2017). While there is a positive probability of deception in any given hand (players need to avoid being "read" too easily by others), on average, individual investments have to reflect a players' assessment of the expected value of his hand. Players who do not invest according to the expected value criterion on average cannot be profitable in the long-run.

Accordingly, highly profitable players are typically described as patient: they have the self-control necessary to stick to a winning strategy by refraining from investing in too many hands. In poker, this concept is known as players' "tightness", and tight players typically invest in no more than 20% of the hands which are dealt to them (Smith et al., 2009). As an illustration of the relevance of the concept of self-control for online poker, Siler (2010) conducted a quantitative analysis of 27 million hands and found the relationship between winning a large

[6] Of course, the purpose of those measures is not to exhaust the notion of skill in the game – the above variables certainly do not – but to point at a number of simple proxies which can be meaningfully manipulated to perform our empirical tests.

proportion of hands and profitability to be *negative* at the player level. This result stresses the necessity for players to be patient and choose the cards they invest in wisely. We therefore define our first proxy for skill in online poker:

Skill 1: self-control. Successful players are *more patient* on average (i.e., they invest in a smaller proportion of their hands).

Further, successful players usually make the most out of the hands in which they invest by playing them aggressively. In poker, this concept is known as players' "aggressiveness", and aggressive players typically raise the amount currently at stake in over 50% of the hands in which they decide to invest (Smith et al., 2009). Such behavior appears highly counter-intuitive for most players, as many believe it to be a good idea to try and "conceal" their lucky draws by refraining from betting too early in the course of a hand (lest that opponents be scared away) (Sklansky and Miller, 2006). Profit maximization requires one to do the contrary. Aggressive betting ensures that the player extracts as much value as possible from the other (dominated) players at the table by building up a large pot right from the start of the hand. To the contrary, not betting aggressively enough allows dominated players at the table to commit less resources at each stage of the game.[7] We therefore define our second proxy for skill in online poker:

Skill 2: aggressive betting. Skillful players are *more aggressive* on average (i.e., conditional on playing a hand, they are more likely to bet aggressively by raising the amount currently at stake).

Finally, the expected value criterion described at the beginning of this section subjects skill 1 and 2 above to players' probability calculation abilities. We therefore define our third proxy for skill in online poker:

Skill 3: probability calculations and hand selection. Successful players (i) *concentrate their investments on stronger hands* (i.e., hands with a higher expected value), and (ii) *bet aggressively on stronger hands* on average.

In the next section, we describe our dataset as well as the variables which we use to operationalize the above skill constructs and conduct our empirical analysis.

3. Data and variables

3.1. Dataset overview

We obtain our data first-hand from Adda52, which is the first, dominant online poker platform in India. The platform was launched in its current form in January 2015. It finances itself by retaining a small, fixed percentage of winning players' earnings in each hand. We receive from Adda52 the raw log files describing the course of action of each hand played at each table from the inception of the site, and until May 2018 (time at which our panel dataset ends). The log files contain the full history of players' activity on the site, whom we can track over time through a unique, constant identifier. For each hand played, the logs notably document players' hole cards (i.e., the two private cards which are dealt to each player at the beginning of each hand), their play decisions at each stage of the hand (e.g., fold, play, bet) and their earnings. We parse this information and, according to a Non-Disclosure Agreement with Adda52, put it in a secured relational database.

We focus our analysis on *No Limit Hold'em ring games*. No Limit Hold'em is the most popular form of poker worldwide. It is most distinctive feature (by contrast to, e.g., Pot Limit Omaha) is that there is no pre-specified limit to how much an individual can bet at any stage of a given hand. Players come to the table with a fixed sum of money called their "stack", which they need to manage. Ring games stand in contrast to tournament structures, where a predetermined number of players from each of several tables moves to the next stage of the competition, until a winner emerges. By contrast, in ring games, each table can be considered independent: players play a succession of poker hands until all lose their money to a single winner or quit the table. The table is then dissolved. We end up with a sample of 91,439 distinct individuals playing No Limit Hold'em ring games over a 3 years and 4 months period. As can be seen from Table 1, our final dataset exhibits a lot of heterogeneity in terms of the total number of hands played by individual players.

Table 1
Distribution of number of hands played across players.

Percentile	Total number of hands
1%	2
5%	2
10%	3
25%	6
50%	28
75%	211
90%	1289
95%	3442
99%	18,140
N players	91,439

Players are randomly matched to playing tables, but have the option to drop-out from a given table at any time (and take their remaining stack money with them). The only information available to players on the platform is the username of their opponents at the table — they have access to no statistics about them. In order to play on the platform, players have to provide an official ID which they need to link to an Indian bank account. This is done to ensure that individuals do not create duplicate accounts on the site.

Players are matched to tables within their selected stake level. The stake level is the minimal *starting* amount a player needs to invest to play his hand — called the "big blind". This amount is typically small, even though bets can escalate very quickly over the course of a hand. At the beginning of each hand, two players are, in turn, forced to put on the table an amount corresponding to half the big blind (called the "small blind"), and the big blind, respectively. This rule ensures that each hand sees some "action", and that players do not always fold their hands.

Within each hand, we only retain the first investment decision of each player at the table, at the exclusion of the two players who are in a forced bet position, and whose investment decisions are not made voluntarily (i.e., the small and big blind bets players). Focusing the analysis on the first round of betting is convenient because all players at the table have to make at least one investment decision irrespective of the quality of the private cards they are dealt. After the first round, some players may have selectively dropped out of the hand, which would need to be accounted for. Players' first investment decision is also simpler to analyze because it is based on their assessment of the strength of their private cards alone, i.e., without regards to the shared community cards which are only revealed in subsequent betting rounds.[8]

Following this strategy, our working dataset includes a total of 85,012,032 player-hand decisions (together with the associated financial outcome) over 40 consecutive months. We structure our data as a panel where one can track player-hand investment decisions and outcomes over time (including the decision to stop playing on the platform). There are 9 stake levels available on the platform, where the big blind amounts to 2, 4, 6, 10, 20, 30, 50, 100 or 2000 rupees

[7] Experienced players also tend to bet aggressively when they bluff, in order to raise opponents' cost of continuing to play and get them to fold before the last stage of the game (where they would have to reveal their cards).

[8] Of course, even in this first round of betting, players who act last benefit from an informational advantage, as they get to observe the first move of others before deciding on their own.

Table 2
Distribution of gains per hand played.

Percentiles	Distrib. of losses (in big blinds)	Distrib. of wins (in big blinds)
1%	−45.25	0.5
5%	−23.6	1.5
10%	−15	1.5
25%	−5.5	2.5
50%	−2	5.5
75%	−1	13
90%	−1	25
95%	−1	34.5
99%	−1	61
N player-hands	15,204,178	27,914,078

The table reads as follows: Among losing hand-player observations, 25% lost at least 5.5 big blinds; alternatively, among winning hand-player observations, 5% won at most 1.5 big blinds. The difference between the number of hand-player observations involved in the construction of this table (15,204,178 + 27,914,078 = 43,118,256) and the total number of observations (85,012,032) correspond to observations for which a player does not invest and whose gains are thus zero.

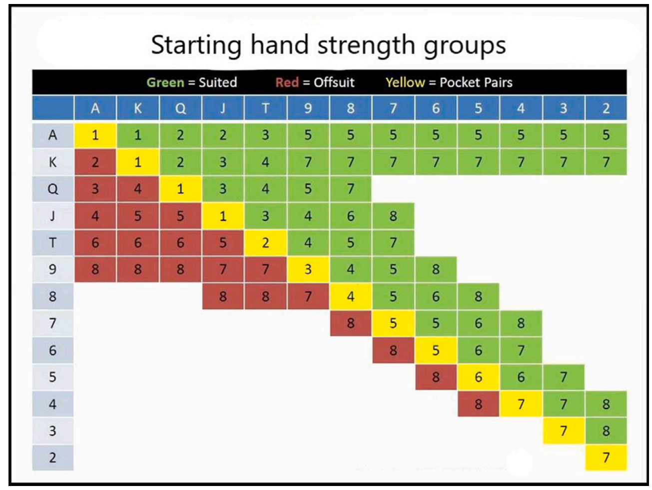

Fig. 1. The expected value or "strength" level of each starting hand.

(exchange rate: 1 rupee ≃ \$0.015). Stake level 2000 is rarely played, and is therefore dropped from the analysis. Throughout this paper, we normalize hand-level earnings by the amount of the big blind so that players' outcomes can be compared across stake levels.

The game of poker typically yields many small wins and losses. Among other characteristics, good players are thus usually described as having the self-control necessary to stick to a winning strategy as they experience short-term, idiosyncratic wins and losses. This important feature of online poker can be seen in our data by computing the distribution of wins and losses (expressed in terms of big blinds) conditional on playing the hand (see Table 2). On average, about 50% of the hands dealt in our data are played (the other half is folded right away and therefore yields a null result). Among those, 75% of the hands lost by players yield a loss of less than 5.5 big blinds. Similarly, 75% of the hands won yield less than 13 big blinds.[9] When played over sufficiently long time-horizons, however, skill in poker typically yields significant profits. As measured by their total earnings, the highest performing decile in our sample earned on average 300 rupees per effective hour spent playing on the platform (i.e., about \$4.6 per hour of play), a sizeable amount for an Indian worker.

3.2. Skill variables

In order to conduct our analysis, we need to operationalize the following constructs:

1. **Self-control:** in poker, players' ability to remain patient and exert the self-control necessary to refrain from investing in too many hands is referred to as their "tightness". It is defined as the average probability of investing some positive amount in a given hand (as opposed to folding it right away). Conversely, players who tend to invest in a significant fraction of the hands dealt to them are called "loose".
2. **Aggressive betting:** conditional on investing in a given hand, players' tendency to bet aggressively is referred to as their "aggressiveness". Aggressive betting often represents a counter-intuitive strategy to follow for novice players. It is defined as the probability that the player raises the amount currently at stake (as opposed to merely matching it).
3. **Hand selection and strength:** in order to measure players' ability to select the hands they invest in according to the expected value criterion, we exploit a unique feature of our data: the fact that it fully documents the private cards which are randomly dealt to all players at the beginning of each hand. From there, we rely on a standard textbook classification of poker hands to classify all hands dealt in 9 exclusive categories of increasing expected value or "strength" (Sklansky and Malmuth, 1999). This classification is reproduced in Fig. 1. We reverse the original coding to get a value of 9 for the best possible hands, and a value of 1 for the worse possible hands.

Based on those constructs, we define our three skill variables of interest from Section 2. We describe those variables at the player level in Table 3. As expected from our above discussion, we can see that the players in our dataset are loose on average, i.e., they have a tendency to invest in a relatively large proportion of the hands dealt to them (about 68%). At the same time, they usually fail to bet aggressively conditional on playing (in 75% of cases). In both situations, the strength of the cards played by players is relatively low on average (1.9 and 3.5 out of a 9-points scale, respectively).

4. Empirical strategy

As a first step to our analysis, we conduct two simple tests of skill predominance in our data by regressing the percentile of earnings achieved by players in each hand played, $Qresult_{p,h}$, on a full set of player fixed effects, γ_p:

$$Qresult_{p,h} = \gamma_p + u_{p,h} \quad (1)$$

This regression explains the variance in hand-level outcomes as a function of inter-personal differences in playing style (captured by the fixed effects). In the case of a pure game of chance, the R^2 from this regression would be statistically indistinguishable from zero: differences in playing style either do not exist, or do not translate into measurable differences in terms of performance. This would for instance be the case for a lottery where players typically have different preferences in terms of the frequency with which they bet on certain numbers, but where those preferences do not impact individual results.

With respect to the legal debate around the dominance of skill in online poker, the criterion understood as "does skill dominate game outcomes" may require the R^2 of this regression to be at least equal to some arbitrary cutoff in order for poker to be considered a game of skill. Understood as "do skilled players dominate the game", the

[9] At the same time, the tails of the distribution indicate that substantial amounts are occasionally won or lost: 1% of losses are larger than 45 big blinds, while 1% of wins are larger than 61 big blinds.

Table 3
Poker decision making variables: Descriptive statistics.

	Skill 1	*Skill 2*	*Skill 3.1*	*Skill 3.2*
	$Pr(play)$ "tightness"	$Pr(raise\|play)$ "aggressiveness"	$E(strength\|play)$ from 1 to 9	$E(strength\|raise)$ from 1 to 9
Mean	0.676	0.253	1.967	3.585
Std	0.233	0.256	1.108	1.778
Observations	91,439	88,910	88,910	70,342

criterion focuses on identifying whether differences in skill between players exist that generate significant profits. This may require the F statistic for the overall significance of this regression to be significant above conventional levels (e.g., 0.1%).

In a second step, we document the impact of our direct measures of skill on individual outcomes. To do so, we regress the average percentile of earnings achieved by players over the hands they played, $Qresult_p$, on our above defined skill variables. Specifically, we test for the impact of self-control and aggressive betting by running the following regression at the player level, where we expect β_1 to be negative and β_2 to be positive:

$$Qresult_p = \beta_0 + \beta_1 play_p + \beta_2 aggress_p + u_p \tag{2}$$

We then introduce hand selection abilities in our model by running the following regression at the player level, where we expect both β_3 and β_4 to be positive:

$$Qresult_p = \beta_0 + \beta_1 play_p + \beta_2 aggress_p + \beta_3 strength|play_p + \beta_4 strength|raise_p + u_p \tag{3}$$

Last, we leverage the detailed panel structure of our data at the player × hand level to estimate the respective contributions of between-player selection and within-player skill acquisition to the population learning process, if any. To test whether learning occurs as a function of experience at the population level, we regress each skill variable in Table 3 on the log of the *current* number of hands played, $n_{ph} = ln(nbhands)_{ph}$. Note, however, that the resulting coefficient of interest on n_{ph} will conflate the learning (within player) and composition effects (between players):

$$skill_{ph} = \beta_0 + \beta_1 n_{ph} + u_{ph} \tag{4}$$

In a second specification, we separate learning from composition effects by explicitly allowing players to select-out of the poker platform (that is, stop playing altogether) as a function of their initial level of skill. To do so, we leverage the fact that our panel tracks *both* the evolution of playing style over time and players' exit decisions. We therefore add the log of the *final* number of hands played by each player: $N_p = ln(final\_nbhands)_p$, so as to effectively clean-up the n_{ph} coefficient of the between-player selection effect:

$$skill_{ph} = \beta_0 + \beta_1 n_{ph} + \beta_2 N_p + u_{ph} \tag{5}$$

The fact that we control for N_p in this second specification implies that the identification of the coefficient on n_{ph} comes from comparing the evolution of skill within strata of players that survive in the platform for the same number of hands. Since the sample is balanced within these strata, our estimate of the evolution of skill as a function of experience will not be biased by selection and will measure pure learning. Conversely, the coefficient on N_p will measure pure selection: because this coefficient is identified keeping n_{ph} fixed, it measures the extent to which players select-out of the platform as a function of their skill *for a given level of experience* n_{ph}. In summary, the coefficients in our model can be interpreted as follows:

- N_p: evolution of skill in the population as a result of selective drop-out (between-player effect),
- n_{ph}: evolution of skill in the population as a result of learning (within-player effect).

Note that this baseline specification does not account for censoring in our data. In other words, because our study covers a fixed period of time (from January 2015 to May 2018), one cannot ensure that the last hand played in our data really is the *final* hand played on the platform. We therefore necessarily identify this final hand, N_p, with some level of noise. In order to assess the robustness of our results to this censoring issue, we re-estimate specification (5) keeping only players whose last hand was played at least m months before the end of our time horizon. This additional constraint aims at ensuring that the last hand played as per our data effectively corresponds to the player's final hand. We then subject the identification of specification (5) to increasingly demanding thresholds for defining player drop-out – from $m = 0$ and up to $m = 6$ months of inactivity by the end of our time period – and report the full distribution of the obtained estimates in Appendix C.

5. Results

5.1. The dominance of skill in online poker

As the first step of our analysis, we rely on regression (1) from Section 4 to provide a simple test of the skill dominance criterion understood as "does skill dominate game outcomes". Our dependent variable of interest, $Qresult_{p,h}$, measures performance according to the percentile of earnings achieved by players in each hand played.[10] This variable thus ranges from 1 (lowest percentile in the distribution of hand-level earnings in the full population of players) to 100 (highest percentile). Regression (1) is run at the player × hand level (with N = 85,111,226) and features the full set of player fixed effects as explanatory variables to account for between-player differences in skill level (not reported). This regression achieves an R^2 of 17.2%. Depending on one's cutoff for skill dominance, this test may therefore lead to the legal conclusion that online poker should be considered a game of chance.

By contrast, within the same empirical framework, a simple test of the "do skill players dominate the game" interpretation of skill dominance yields unambiguous support to the hypothesis that skill drives performance at the player level. The F statistic for the overall significance of this regression yields F = 5.98 and $p < 0.001$. To visualize this more clearly, consider the 90th percentile of players in our data in terms of the total number of hands played on the platform. This population represents a balanced sample of 9,145 players who played (at least) 1,289 hands in total (see Table 1). Within this balanced sample, Fig. 2 reports the significance of the F statistic as increasing information on players' playing history – proxied by their current number of hands played, from 2 to 1,289 – is added to regression (1). We can see that the p-value associated with this significance test drops sharply as we allow the model to access more information on hand history within our sample of players. Specifically, the model is able to reject the hypothesis that skill does not drive player performance in this sample with a probability of $p < 0.001$ with as little information as 7 hands played (corresponding to slightly over 5 min of effective play in our data).

In a second step, we provide a direct test of the impact of skill on player-level performance. We use our simple proxies for skill for

[10] Remember that we normalize hand-level results by the amount of the big blind to ensure comparability across stake levels.

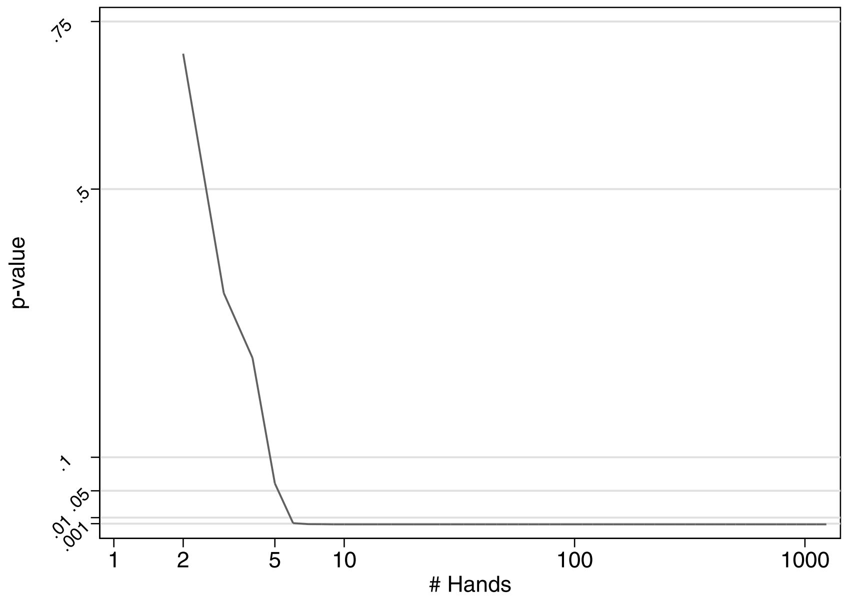

Fig. 2. The significance of skill as a function of considered play history. The graph depicts the overall significance (F statistic) of regression (1) as a function of the length of the hand-playing history considered in the model (from 2 to 1289 hands played). The sample of players considered is fixed across regressions (N = 9145), and corresponds to the 90th percentile of the distribution in terms of the total number of hands played on the platform.

this purpose, and report the results of regressions (2) and (3) from Section 4 in columns (1) and (2) of Table 4. In column (1), we regress our indicator of financial performance averaged at the player level, $Qresult_p$, on our proxies for self-control and aggressive betting. The model has therefore two explanatory variables: (i) the average probability to play a given hand ("tightness"), and (ii) the average probability to raise conditional on playing the hand ("aggressiveness"). We can see that loose players achieve significantly lower performance on average, while the reverse holds for aggressive players. This result confirms that a restrictive application of the skill dominance criterion understood as "does skill dominate game outcomes" can conceal the fact that, for a minority of elite players, skill does generate sustained profits.

In column (2), we further test for the impact of players' ability to select hands with a higher expected value (or strength), conditional on both playing (i.e., investing in the hand) and raising (i.e., betting aggressively). Consistent with our expectations, we can see that the coefficients on both strength variables are positive and significant. That is, for a given level of tightness and aggressiveness at the player level, players who demonstrate an ability to select stronger hands on average achieve better results. All in all, our empirical results provide unambiguous support for the idea that skill drives performance and earnings in online poker (i.e., "skilled players dominate the game").

In Appendix A, we discuss the possible mediating effect of the skill level of the player's opponents on our estimates. We show that, if anything, the level of skill of the opponents at the table attenuates the relationship between own-skill and performance. In other words, strong opponents are able to reduce the amount of profit that can be reaped from skillful play. In Appendix B, we explore the sensitivity of our tightness and aggressiveness coefficients to excluding players who exceed the approximate target values for average optimal play of 20% tightness (4.2% of players in our data), and 50% aggressiveness (14.5% of players) which we derive in Section 2. The analysis demonstrates that excluding the minority of players who exceed those benchmarks for skillful play delivers the same empirical conclusion as to the strength of the relationship between skill and performance (although with increased estimated effect sizes).

Table 4
The impact of skill on player performance.

	(1)	(2)
	$Qresult_p$	$Qresult_p$
$Pr(play)_p$ [*"tightness"*]	−8.269***	−3.186***
	(0.260)	(0.367)
$Pr(raise\|play)_p$ [*"aggressiveness"*]	6.170***	0.843**
	(0.280)	(0.333)
$E(strength\|play)_p$ [*hand selection*]		0.440***
		(0.114)
$E(strength\|raise)_p$ [*hand selection*]		0.112***
		(0.0300)
N	88,910	70,342
R^2	0.035	0.009

Regression type: cross-sectional, player level. The table presents OLS estimates with robust standard errors clustered at the player level in parentheses (constant not reported). ***, ** and * denote statistical significance at the $p < 0.001$, $p < 0.01$ and $p < 0.05$ levels, respectively.

5.2. The evolution of skill: Learn or Leave?

In this section, we proceed to test for learning in the underlying population of players. To do so, we start by relating each ability variable (as defined in Table 3) to players' level of experience (measured in rank percentile of the number of hands already played on the platform) in Fig. 3. The picture provides a non parametric representation of the evolution of playing style with experience, together with the 95% confidence interval around each mean value (although those are too small to be visible in the graph).

Starting from the top-left figure, we can see that the probability to invest some positive amount in any given hand (i.e., "looseness") decreases from almost 70% among beginners to about 25% among most experienced players – a number which approaches the reference value of 20% often cited for tight poker play. Relatedly, the average strength of the hands played increases with experience, from a value of less than 2 (out of 9) to a value of 4. Finally, the probability that players bet aggressively (i.e., raise the amount currently at stake conditional on investing) increases from about 20% among beginners to 60% among most experienced players – a threefold increase in likelihood – to reach levels which are more in line with expert play.

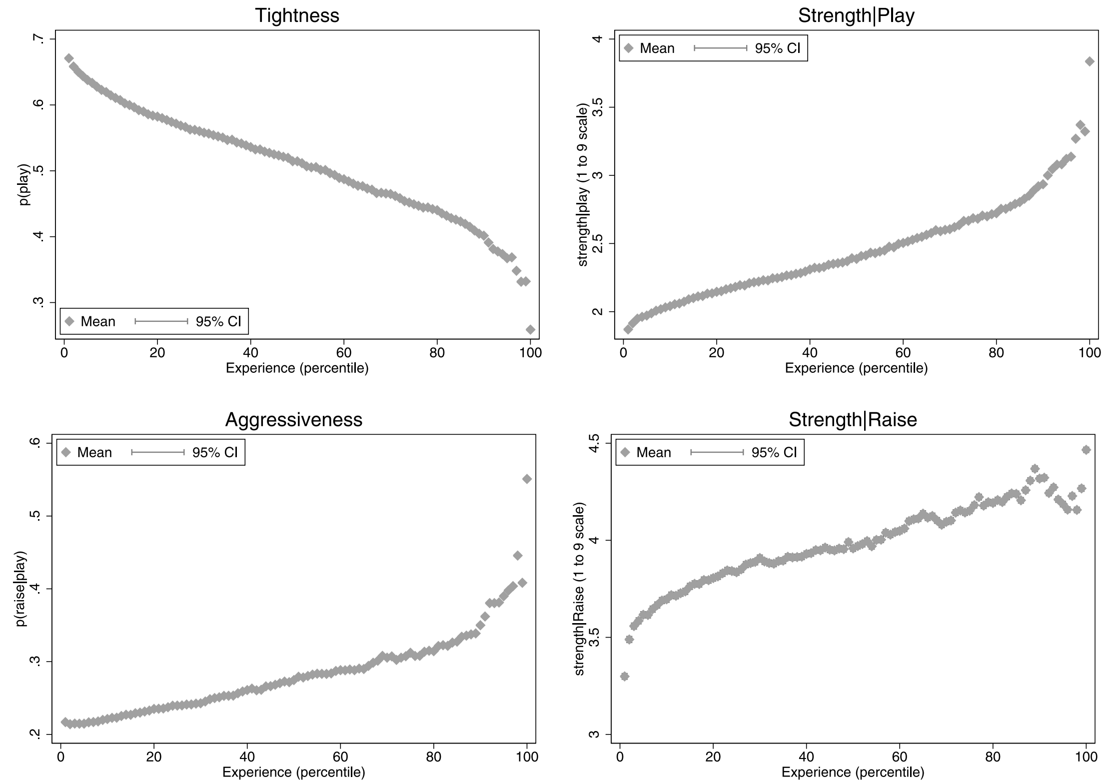

Fig. 3. The evolution of skill with experience.

Overall, we find strong evidence of learning across all four skill measures at the population level. It remains unclear, however, how much of this learning process occurs as a result of players' selective drop-out from the platform as opposed to skill acquisition. As the last step of or our analysis, we proceed to study the extent to which players (i) learn and acquire playing skills as a function of relevant field experience, and (ii) selectively drop-out of the poker platform as a function of initial ability. That is, we seek to establish the extent to which less able players *learn* or *leave*.

To this effect, we report the results of regressions (4) and (5) from Section 4 in Table 5. Columns (1)–(4) confirm our above findings that learning is significant at the population level. Each percentage increase in the number of hands played is associated with a 3.9% decrease in looseness (column (1)), a 2.2% increase in aggressiveness (column (3)), and a 0.15 and 0.10 increase in strength conditional on playing or betting, respectively. Columns (5)–(8) decompose the population learning coefficients in two parts: between-player selection (N_p), and within-player learning (n_{ph}). Across columns (5)–(7), between-player selection and within-player learning appear to contribute roughly equally to the population learning process. The only exception is column (8) ($E(strength|raise)$), where learning seems to occur exclusively between players (i.e., as a result of selective drop-out).

In Appendix C, we assess the robustness of those "Learn of Leave" results by reporting the distribution of our point-estimates under a number of additional censoring constraints aimed at more reliably identifying player drop-out in our data. Specifically, we limit our estimating sample to players that remain inactive for (at least) the last 30 days covered by our study — and up to 6 months. For this subset of players, it is more likely that the last hand played on the platform corresponds to their final hand, such that N_p is measured more precisely. The coefficients we obtain are similar and all fall within the same margin of error.

In Fig. 4, we rely on our results from columns (5)–(7) in Table 5 to provide a visual representation of the estimated player skill acquisition process (i.e., "within-player learning") as a function of their initial level of ability (in percentile of the distribution of skill in the population). More precisely, for each starting level of skill (x-axis), the figure reports the estimated amount of practice necessary to reach the skill level exhibited by the top decile of players in terms of experience. For interpretability, we report players' required level of practice on the y-axis in terms of both (i) the number of hands played, and (ii) the effective time spent playing on the platform.[11]

We can see that, depending on their initial skill, players differ widely in the amount of practice necessary to reach the ability level of the most experienced players in the platform. All in all, for a novice player starting at the bottom decile of the distribution of skill, our estimates indicate that reaching the skill level of the most experienced players requires one to play about 57,000 hands on average, corresponding to 710 h of practice (i.e., 30 days). For an individual aiming at an average of 5 h of effective time spent playing per working day, this corresponds to about 7 months of full-time training. Of course, this number should not be taken too literally, as our skill variables remain relatively simple and do not exhaust the notion of skill in online poker. In this sense, our estimate might be interpreted as a lower bound on the time required to become a skilled player.

6. Legal implications

Similar to e-sports and other "sportified" games of skill (Hallmann and Giel, 2018; Heere, 2018; Meng-Lewis et al., 2021), the practice of online poker has exploded following the advent of the Internet. In this context, courts typically approached the question of whether online poker could be legally distinguished from gambling (i.e., a game of chance), based on an evaluation of the predominance of the element of skill in the game. So far, however, the judge has operated in a

[11] As can be seen from the graph, the average hand lasts about 45 s in our data.

Table 5
Learn of leave?

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$Pr(play)$	$E(strength\|play)$	$Pr(raise\|play)$	$E(strength\|raise)$	$Pr(play)$	$E(strength\|play)$	$Pr(raise\|play)$	$E(strength\|raise)$
n_{ph}	−0.039***	0.152***	0.022***	0.103***	−0.019***	0.079***	0.018***	−0.002
	(0.002)	(0.007)	(0.002)	(0.011)	(0.001)	(0.004)	(0.001)	(0.005)
N_p					−0.027***	0.096***	0.005***	0.136***
					(0.002)	(0.007)	(0.002)	(0.014)
N	85,111,226	43,118,084	43,118,084	11,841,897	85,111,226	43,118,084	43,118,084	11,841,897
R^2	0.026	0.016	0.010	0.006	0.028	0.018	0.010	0.008

Regression type: panel, player × hand level. The table presents OLS estimates with robust standard errors clustered at the player level in parentheses (constant not reported). ***, ** and * denote statistical significance at the $p < 0.001$, $p < 0.01$ and $p < 0.05$ levels, respectively.

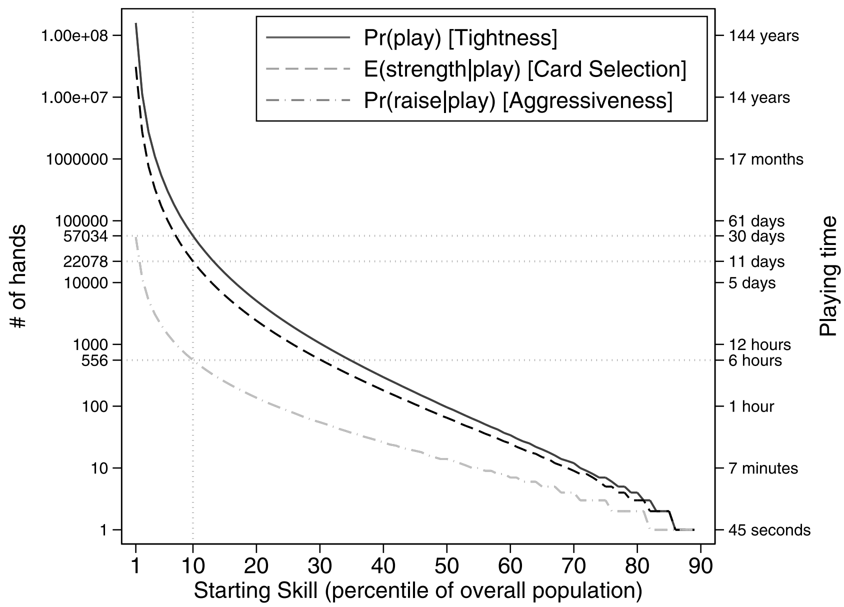

Fig. 4. Practice Required to Reach the Skill Level of Highly Experienced Players (top decile). **Note:** for different starting levels of skill (x-axis), the figure reports the amount of practice required to match the ability of the most experienced players on the platform (90-th percentile of hands played). Practice is expressed in terms of (i) number of hands played (left y-axis), and (ii) playing time (right y-axis). Playing time is computed for a hand duration of 45 s (sample average). Starting skill levels are expressed in percentile of the population distribution of skills. The figure reads as follows: a player starting in the bottom decile of skill in terms of tightness would need to play at least 57,034 hands (corresponding to about 30 days or 710 h of training) to reach the tightness level of the most experienced players (top decile). The variable $E(strength|raise)$ is omitted from the graph as the related coefficient of interest on n_{ph} (Table 5, column (8)) fails to reach statistical significance.

methodological vacuum that may precisely result from the ambivalence of the notion of skill dominance (Miles et al., 2013; Sayta, 2012).

In this paper, we argue that the concept of skill dominance might be better understood as “do skilled players dominate the game” than as “does skill dominate game outcomes”. The reason is simple: the latter criterion puts uncertain decision making environments at a (questionable) disadvantage in the legal evaluation process. Because many decision making fields of interest to law and economics are inherently uncertain (e.g., finance or management), we argue that the former approach is more commonly used to characterize skill across occupations, as it merely requires that differences in skill exist between players that generate sustained profits (Bertrand and Mullainathan, 2001; Barber and Odean, 2000, 2013; Akepanidtaworn et al., 2022).

We derive the empirical consequences of the above legal distinction by proposing a simple regression framework in which both interpretations can be tested independently. Depending on the threshold retained, our skill dominance tests may indeed lead to opposite legal conclusions. Skill in our population of players “only” explains about 17.2% of the variance in hand-level outcomes. At the same time, our model is able to reject the hypothesis that skill does not drive player performance with as little as 7 hands of play history ($p < 0.001$).

The second contribution of our paper is to advance our understanding of both the nature of skill dominance in online poker, and of the learning ecosystem that underpins the game. We provide simple but direct proxies for skill in the game that demonstrate how skillful play impacts player performance. Specifically, we single out: (i) self-control, (ii) an ability to take (calculated) risk, and (iii) good probability calculation and hand selection.

From there, we develop an empirical procedure that allows us to describe: (i) whether less able players quit playing the game as a function of starting ability (i.e., the exit rate), and (ii) whether more able players invest in training and improve their skills (i.e., the learning rate). Both statistics are useful to the analyst, as they are indicative of a well-functioning competitive ecosystem: less able players learn about their ability and leave the platform, while more able players invest in training and learn.

We conclude from our conceptual and empirical analyses that the scholarly debate around the practice of online poker may move beyond that of its legality, and focus instead on deriving sound regulating and choice architecture principles for this industry — both at the recreational and competitive levels. Such principles might for instance include the systematic provision of targeted information to newcomers, or the establishment of (possibly customizeable) commitment rules and alerts aimed at improving players’ bankroll management (see, e.g., Heimer and Imas, 2022 for a recent application in the context of stock trading). We leave these regulation and choice architecture questions open for future research.

Beyond the particular case of online poker, the conceptual and empirical treatment of the legal notion of skill dominance which we present in this paper may apply to any (presumed) game of skill. As long as the activity in question involves money management, similar issues of legality may arise before courts, and the judge may have to decide on whether the activity under consideration constitutes online gambling.

Sports betting emerges as an obvious example. Over the past decades, the (illegal) practice of betting on sports events has risen sharply in the U.S., generating between $80 and $380 billion a year in profits for offshore websites (and organized crime).[12] On May 14, 2018, the Supreme Court of the United States struck down the 1992 Professional and Amateur Sports Protection Act (PASPA) for unconstitutionality, arguing that the federal prohibition on sports betting was infringing on state legislative bodies (Holden, 2018). This decision has paved the way for a broad movement of legalization of sports betting across federal states. However, domestic players in this industry continue to operate in a gray zone from a legal perspective. This is especially true when it comes to the 2006 Unlawful Internet Gambling Enforcement Act (UIGEA), as the question of whether sports betting constitutes online gambling has not yet arisen before courts.

Our paper provides a number of tools and routines that can be used to assess the skill dominance criterion in such contexts. Our analysis only requires (i) a panel data set where players' decisions and outcomes can be tracked over time, and (ii) the definition by the game's practitioners of a few meaningful proxies for skill that can be used for analysis. From there, our empirical procedures can be easily replicated to evaluate the skill dominance criterion from a legal standpoint (directly or indirectly), as well as characterize the quality of the game's competitive ecosystem in terms of (i) players' learning rate, and (ii) their drop-out rate as a function of ability.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors do not have permission to share data. The data may be obtained upon request directly from Adda52.

Appendix A. The impact of opponent skill on player performance

In this Appendix, we seek to estimate whether the impact of own skill on player-level performance depends on the skill level of his or her best opponent at the table. In order to achieve this goal, we augment Table 4 in the main text by (i) controlling for the skill level of the best opponent at the table (approximated by their current level of experience as per the log-number of hands they have played so far), and (ii) interacting the effect of the skills of the player under consideration with the level of experience of his or her best opponent at the table.

In other words, we exploit the fact that players are randomly matched within a given stake level while differing in their current level of experience to identify how the level of skill of the best opponent at the table mediates the relationship between own skill and performance. Note, however, that players select the stake level at which they want to play, which will induce a positive correlation between own and opponent skill. In order to get the intuition for this, one can turn to Fig. A.1, which relates the average effective time spent playing on the platform (i.e., experience) to players' most played stake level (as well as Table A.1 which describes the proportion of hands played at each stake level in our data). We can see from Fig. A.1 that the experience level of players tends to co-evolve across stake levels. In order to break this correlation, we estimate our model at the player × stake level and control for stake level fixed effects in our regressions.

Table A.1
Proportion of hands played at each stake level.

Stake level	Amount in rupees	Prop. hands played
1	2	0.14
2	4	0.23
3	6	0.06
4	10	0.12
5	20	0.17
6	30	<0.01
7	50	0.15
8	100	0.14
9	2000	<0.01

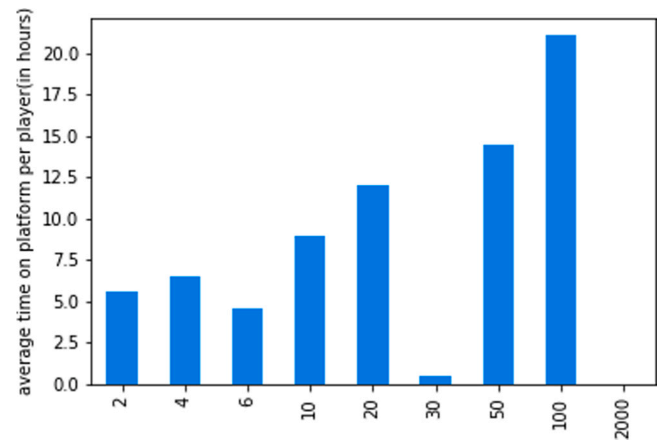

Fig. A.1. Player experience as a function of most played stake level. The graph assigns each player in the data to the stake level (reported in rupees) in which they have played the highest proportion of their hands. It then reports the average effective time spent playing on the platform across players as a function of their most played stake level.

The results are reported in Table A.2. When we do not control for stake level fixed effects (columns (1) and (2)), we find a positive coefficient of opponent's skill. By contrast, the coefficient turns negative once we control for stake levels (columns (3) and (4)). This confirms the fact that own and opponent skill are positively correlated across stake levels, hence the positive bias on the OLS coefficient in the absence of stake level controls.

Turning our attention to columns (5) and (6), we find mixed empirical support for the proposition that the level of experience of the best opponent at the table mediates the impact of own skill on performance. This can be seen by looking at the coefficients on the interaction terms, which are mostly statistically insignificant. The only exception relates to the variable $Opponent\_Skill \times E(strength|play)$, which features a negative coefficient. This implies that players' ability to better select the hands in which they invest yields smaller payoffs against more experienced opponents. In other words, strong opponents are able to significantly reduce the amount of profit that can be reaped from skillful play along this dimension.

Appendix B. Skill and performance: Sensitivity analysis

In Section 2, we derive rough theoretical target values for what may constitute average optimal play in terms of player tightness (i.e., $Pr(play)_p \approx 20\%$), and aggressiveness (i.e., $Pr(raise|play)_p \approx 50\%$). In this Appendix, we analyze the sensitivity of the corresponding coefficients (reported in column (1) of Table 4 in the main text) to excluding the 4.2% of players who exceed the 20% threshold for average tightness in our data, and the 14.5% of players who exceed the 50% threshold for average aggressiveness, respectively.

[12] See the National Gambling Impact Study Commission Final Report: http://govinfo.library.unt.edu/ngisc/reports/2.pdf.

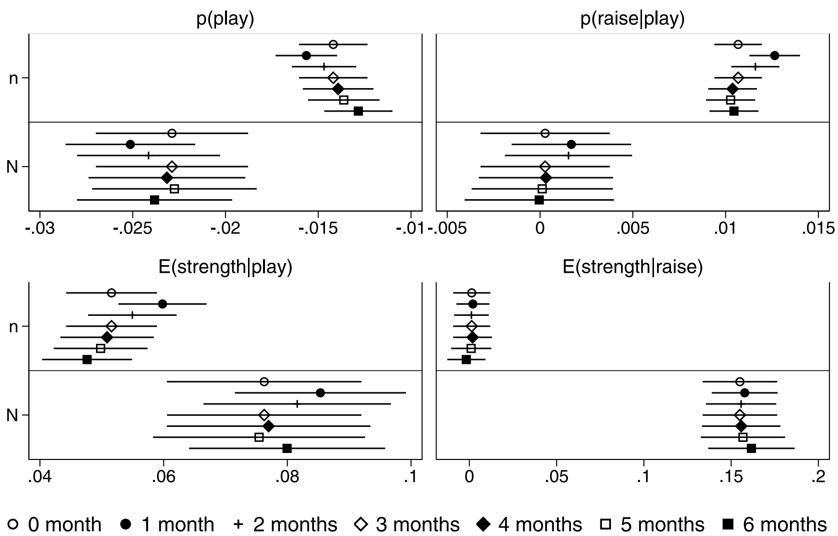

Fig. A.2. The effect of censoring on the "*Learn or Leave*" Estimates. **Note:** for each skill variable in Table 5, this figure reports the distribution of the N_p and n_{ph} "*Learn or Leave*" coefficients obtained from specification (5) in Section 4, subject to imposing an increasingly demanding threshold for defining player drop-out in the data — from zero and up to 6 months of inactivity by the end of our time period.

Table A.2
The impact of opponents' skill level on player performance.

	(1) $Qresult_{p,stake}$	(2) $Qresult_{p,stake}$	(3) $Qresult_{p,stake}$	(4) $Qresult_{p,stake}$	(5) $Qresult_{p,stake}$	(6) $Qresult_{p,stake}$
$Opponent\_Skill_{p,stake}$	0.105***	0.124***	−0.142***	−0.168***	−0.147***	−0.158***
	(0.0276)	(0.0300)	(0.0309)	(0.0346)	(0.0249)	(0.0233)
$Pr(play)_{p,stake}$	−5.971***	−1.627***	−4.836***	−0.686***	−4.831***	−0.782***
	(0.156)	(0.235)	(0.161)	(0.235)	(0.160)	(0.235)
$Opponent\_Skill \times Pr(play)$					0.0752	−0.0470
					(0.132)	(0.186)
$Pr(raise\|play)_{p,stake}$	6.873***	2.522***	6.867***	2.571***	6.867***	2.604***
	(0.166)	(0.202)	(0.165)	(0.201)	(0.164)	(0.200)
$Opponent\_Skill \times Pr(raise\|play)$					−0.0182	0.0287
					(0.130)	(0.166)
$E(strength\|play)_{p,stake}$		0.407***		0.322***		0.298***
		(0.0660)		(0.0661)		(0.0657)
$Opponent\_Skill \times E(strength\|play)$						−0.102**
						(0.0509)
$E(strength\|raise)_{p,stake}$		0.120***		0.0895***		0.0899***
		(0.0180)		(0.0180)		(0.0179)
$Opponent\_Skill \times E(strength\|raise)$						−0.0197
						(0.0174)
Stake level fixed effects	NO	NO	YES	YES	YES	YES
N	221,321	176,836	221,321	176,836	221,321	176,836
R^2	0.036	0.010	0.047	0.023	0.047	0.024

Regressions are at the player × stake level. The table presents OLS estimates with robust standard errors clustered at the player level in parentheses (constant not reported). Columns (1) and (2) do not control for stake level fixed effects, therefore inducing a positive correlation between own and opponent skill. Columns (3)–(6) break this correlation by including stake level fixed effects. ***, ** and * denote statistical significance at the $p < 0.001$, $p < 0.01$ and $p < 0.05$ levels, respectively.

We report our results in Table A.3. We can see from column (1) that excluding players who exceed our benchmark for tightness from the analysis delivers a relatively higher, yet similar coefficient. This is indicative of the fact that our target value for optimal tightness might be slightly overestimated: on average, players may move below the 20% target and continue to reap additional benefits in terms of increased performance. By contrast, we can see from column (2) that excluding players who exceed our benchmark for aggressiveness results in a significant increase in the estimated effect of aggressive play on individual earnings: from 6.17 percentiles in Table 4, to 15.25 percentiles in Table A.3. This suggests that the relationship between aggressive play and performance changes significantly once players exceed our derived target value, so that they may not move significantly above this threshold without experiencing a decrease in relative performance.

Overall, however, our main result that skill significantly drives player performance remains true irrespective of whether we estimate the relationship between skill and individual earnings in the full population of players (i.e., including players who exceed our benchmarks), or in the sub-population of players who are more directly concerned by our theoretical predictions (i.e., retaining only players who do not exceed those benchmarks).

Appendix C. Effect of censoring on "*Learn or Leave*" estimates

Our baseline "*Learn or Leave*" analysis from Section 5.2 does not explicitly account for censoring in our data. Because our study covers a fixed period of time, however, one cannot ensure that the last hand played by our players on the platform is, in fact, their final hand. This necessarily introduces noise in our identification of player drop-out which, if correlated with unobservables, may have an impact on our estimates. In order to address this issue, this Appendix imposes a number of additional constraints aimed at better identifying player drop-out in the data.

Table A.3
The impact of skill on player performance: Sensitivity analysis.

		(1)	(2)
		$Qresult_p$	$Qresult_p$
$Pr(play)_p$	[*"tightness"*]	−8.588***	−7.075***
		(0.265)	(0.250)
$Pr(raise\|play)_p$	[*"aggressiveness"*]	6.252***	15.25***
		(0.283)	(0.315)
Excluded players		"too tight"	"too aggressive"
N		88,525	78,160
R^2		0.036	0.067

Regression type: cross-sectional, player level. The table presents OLS estimates with robust standard errors clustered at the player level in parentheses (constant not reported). ***, ** and * denote statistical significance at the $p < 0.001$, $p < 0.01$ and $p < 0.05$ levels, respectively.

Specifically, to consider a player's last hand as their final hand played on the platform (i.e., the player has effectively selected-out), we require that they have been inactive for (at least) m full months by the end of our time period — where m varies from 0 and up to 6 consecutive months. We then remove players whose last hand does not satisfy this constraint from the estimation and compare the resulting "*Learn or Leave*" estimates to those reported in the main text, which are based on the full sample of players. In Fig. A.2, we report the full distribution of these "*Learn or Leave*" estimates. We can see that, irrespective of the censoring constraints imposed on the data, our coefficients of interest on N_p (between-player selection) and n_{ph} (within-player learning) all remain of the same magnitude and fall within the same margin of error.

References

Akepanidtaworn, K., Di Mascio, R., Imas, A., Schmidt, L., 2022. Selling fast and buying slow: Heuristics and trading performance of institutional investors. J. Finance.

Barber, B.M., Odean, T., 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. J. Finance 55 (2), 773–806.

Barber, B.M., Odean, T., 2013. The behavior of individual investors. In: Handbook of the Economics of Finance, Vol. 2. Elsevier, pp. 1533–1570.

Bertrand, M., Mullainathan, S., 2001. Are ceos rewarded for luck? the ones without principals are. Q. J. Econ. 116 (3), 901–932.

Bowling, M., Burch, N., Johanson, M., Tammelin, O., 2015. Heads-up limit hold'em poker is solved. Science 347 (6218), 145–149.

Croson, R., Fishman, P., Pope, D.G., 2008. Poker superstars: Skill or luck? similarities between golf–thought to be a game of skill–and poker. Chance 21 (4), 25–28.

DeDonno, M.A., Detterman, D.K., 2008. Poker is a skill. Gaming Law Rev. 12 (1), 31–36.

Duersch, P., Lambrecht, M., Oechssler, J., 2020. Measuring skill and chance in games. Eur. Econ. Rev. 127, 103472.

Gabaix, X., Landier, A., 2008. Why has ceo pay increased so much? Q. J. Econ. 123 (1), 49–100.

Hallmann, K., Giel, T., 2018. Esports–competitive sports or recreational activity? Sport Manag. Rev. 21 (1), 14–20.

Heere, B., 2018. Embracing the sportification of society: Defining e-sports through a polymorphic view on sport. Sport Manag. Rev. 21 (1), 21–24.

Heimer, R.Z., Imas, A., 2022. Biased by choice: How financial constraints can reduce financial mistakes. Rev. Financ. Stud. 35 (4), 1643–1681.

Holden, J.T., 2018. Prohibitive failure: The demise of the ban on sports betting. Georgia State Univ. Law Rev. 35, 329.

Levitt, S.D., Miles, T.J., 2014. The role of skill versus luck in poker evidence from the world series of poker. J. Sports Econ. 15 (1), 31–44.

McCormack, A., Griffiths, M.D., 2012. What differentiates professional poker players from recreational poker players? a qualitative interview study. Int. J. Mental Health Addict. 10 (2), 243–257.

Meng-Lewis, Y., Wong, D., Zhao, Y., Lewis, G., 2021. Understanding complexity and dynamics in the career development of esports athletes. Sport Manag. Rev..

Meyer, G., von Meduna, M., Brosowski, T., Hayer, T., 2013. Is poker a game of skill or chance? a quasi-experimental study. J. Gambl. Stud. 29 (3), 535–550.

Miles, T.J., Levitt, S., Rosenfield, A.M., 2013. Is texas hold'em a game of chance? a legal and economic analysis. Georget. Law J. 101, 581.

Moravčík, M., Schmid, M., Burch, N., Lisỳ, V., Morrill, D., Bard, N., Davis, T., Waugh, K., Johanson, M., Bowling, M., 2017. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. Science 356 (6337), 508–513.

Sayta, J., 2012. Legality of poker and other games of skill: A critical analysis of india's gaming laws. NUJS Law Rev. 5, 93.

Siler, K., 2010. Social and psychological challenges of poker. J. Gambl. Stud. 26 (3), 401–420.

Sklansky, D., Malmuth, M., 1999. Hold'em Poker for Advanced Players. Two Plus Two Publishing LLC.

Sklansky, D., Miller, E., 2006. No Limit Hold'em: Theory and Practice. Two Plus Two Publishing LLC.

Smith, G., Levere, M., Kurtzman, R., 2009. Poker player behavior after big wins and big losses. Manage. Sci. 55 (9), 1547–1555.

Spapens, T., 2014. Illegal gambling. In: The Oxford Handbook of Organized Crime. pp. 402–418.

Van Essen, M., Wooders, J., 2015. Blind stealing: Experience and expertise in a mixed-strategy poker experiment. Games Econom. Behav. 91, 186–206.

van Loon, R.J.P., van den Assem, M.J., van Dolder, D., 2015. Beyond chance? the persistence of performance in online poker. PLoS One 10 (3).