Helsinki GSE

HELSINKI GSE DISCUSSION PAPERS 10 · 2023

Skill-Related Differences in Reference-Dependent Stopping

Marco Lambrecht

HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

HANKEN

Aalto University

Helsinki GSE Discussion Papers

Helsinki GSE Discussion Papers 10 · 2023

Marco Lambrecht:
Skill-Related Differences in Reference-Dependent Stopping

ISBN 978-952-7543-09-2 (PDF)
ISSN 2954-1492

Helsinki GSE Discussion Papers:
https://www.helsinkigse.fi/discussion-papers

Helsinki Graduate School of Economics
PO BOX 21210
FI-00076 AALTO
FINLAND

Helsinki, Revised November 2025

Skill-Related Differences in Reference-Dependent Stopping

Marco Lambrecht*

November 10, 2025

Abstract

This paper examines reference-dependent stopping in repeated risky decision-making and its relation to skill dependence and individual ability. Using field data from online heads-up poker tournaments, we classify players according to whether they are more likely to stop when ahead (gain-exit behavior) or when behind (stop-loss behavior). Skilled players, relative to less skilled players, are more likely to adopt gain-exit behavior and play more frequently, yet their performance deteriorates during periods of chasing losses. These patterns suggest that skilled individuals may pursue implicit income targets, continuing play until reaching them, which can undermine short-term performance and unintentionally generate left-skewed profit distributions. The findings raise new questions about how skill and perceived control shape behavior in repeated risk-taking, offering insights for other domains such as trading and performance-based work.

Keywords: reference dependence, stopping behavior, repeated risk-taking, skill, poker

JEL-Codes: D81, D91, G41

---

*University of Stavanger and Hanken School of Economics & Helsinki GSE
Email: marco.lambrecht@uis.no
Address: UiS School of Business and Law, N-4036 Stavanger, Norway

I like to thank Mark Dean for valuable discussions within the "Adopt a Paper"-program managed by Daniel Gomez, as well as Aleksandr Alekseev, Benjamin Bushong, Grigory Chernov, Peter Duersch, Sebastian Ebert, Katrin Gödker, Daniel Hauser, Jamie Lien, Topi Miettinen, Jörg Oechssler, Jussi Palomäki, Illia Pasichnichenko, Hannes Rau, Andis Sofianos, Niko Suhonen, Agnieszka Tymula, Michael Ungeheuer, Christoph Vanberg, Martin Weber and the seminar audiences at University of Heidelberg, Helsinki GSE, and FUR Ghent for their helpful comments. Financial support by DFG grant OE 198/5-1 is gratefully acknowledged.

1 Introduction

Repeated risk-taking is common across many domains of economic life. From financial trading and speculative investing to professional gaming and sports betting, individuals frequently make a series of risky decisions where they must determine not only how much to risk, but also whether to continue. Such environments often give rise to systematic behavioral patterns, most notably the tendency to chase losses—i.e., persisting in risk-taking after setbacks in an attempt to recover. In financial markets, this behavior has been linked to extreme episodes of rogue trading, where traders escalate positions after losses with disastrous consequences for their institutions, as in the cases of Nick Leeson, Kweku Adoboli, and Jérôme Kerviel (Fraser-Mackenzie et al., 2019). In most real-world contexts—including trading, entrepreneurship, and games like poker—outcomes are not governed purely by chance, but depend jointly on luck and skill. What remains less well understood is how the skill dependence of an environment, and individuals' own skill levels, shape the emergence and persistence of such reference-dependent patterns.

This paper examines how skill influences reference-dependent stopping behavior in repeated risky decision-making. Reference-dependent preferences are well established in the study of decisions under risk (e.g., Kőszegi and Rabin, 2006), and two commonly observed tendencies are the house money effect and chasing behavior. The house money effect describes the tendency to take on greater risk when ahead of one's reference point (Thaler and Johnson, 1990), whereas chasing behavior refers to increased risk-taking when behind, motivated by the desire to "break even" (Lesieur, 1977). From a behavioral perspective, these reference-dependent patterns in risk-taking have direct analogues in stopping behavior. The house money effect implies a form of stop-loss strategy, whereby decision makers tend to continue when ahead and are more inclined to stop after losses. Conversely, chasing losses resembles a gain-exit strategy, in which individuals are less likely to stop when behind but more likely to stop after gains.

Although these patterns have been observed in numerous experimental and field settings, little is known about how they vary across environments that differ in their dependence on skill or across individuals of differing ability. We address this question by analyzing decisions of online poker players in a setting that provides a rare combination of repeated, high-stakes choices and precise measurement of both luck and skill. Specifically, we exploit data from heads-up "Sit-and-Go" poker tournaments, a format that isolates stopping decisions while holding strategic and institutional features largely constant. Because each match ends deterministically when one player wins all chips and a new match begins against a different opponent, the choice of whether to continue playing after a gain or loss provides a clean field measure of reference-dependent stopping.

To quantify the role of skill, we use the methodology of Duersch et al. (2020) to estimate individual ability levels and the overall skill dependence of each poker environment. We then classify players according to whether they are more likely to stop playing when ahead (gain-exit) or when behind (stop-loss), and test how these stopping patterns relate to differences in skill, stakes, and game speed. Variation in game speed allows us to exploit systematic differences in the degree to which outcomes depend on skill rather than luck, while differences in stakes enable us to examine whether the observed behavioral patterns are robust across monetary incentive levels. We also analyze how behavioral types differ across several performance dimensions, including frequency of play, overall profitability, performance conditional on being ahead or behind, the skewness of session profits, and the timing of their next session. This design allows us to connect stopping patterns to differences in skill and environment, and to explore how these patterns relate to broader measures of behavior and performance, extending prior work in which these connections have remained largely unexplored..

An extensive empirical literature documents reference-dependent stopping patterns across economically relevant domains. In financial markets, investors and traders adjust continuation and risk after gains and losses (e.g., Barberis et al., 2001; Coval and Shumway, 2005; Zhang and Semmler, 2009; Huang and Chan, 2014). Similar patterns arise in gambling environments, including horse betting (Suhonen and Saastamoinen, 2018) and online poker (Smith et al., 2009; Eil and Lien, 2014). To probe the mechanisms behind such behaviors, recent studies distinguish between realized and "on paper" outcomes, showing that individuals are more likely to continue risk-taking after paper losses and unrealized gains (Imas, 2016; Merkle et al., 2021). When it comes to individual-level determinants, other research has examined how experience and past success relate to behavioral biases (e.g., List, 2003; Feng and Seasholes, 2005; Locke and Mann, 2005; Fraser-Mackenzie et al., 2019). Yet no study has systematically investigated how measured individual skill and the skill dependence of the environment are related to reference-dependent stopping behavior.

A complementary line of research studies *income targeting* in dynamic work and earning decisions. In the classic day-labor setting of New York City taxi drivers, Camerer et al. (1997) show that drivers are more likely to stop after meeting daily earnings targets, a result that sparked a rich follow-up literature refining measurement and modeling of reference points in intertemporal labor supply (e.g., Farber, 2008; Crawford and Meng, 2011; Thakral and Tô, 2021). These studies focus on whether people stop earlier after good outcomes and continue working after shortfalls. The present paper contributes to this literature by examining stopping in a strategic, competitive environment where outcomes reflect both luck and skill. This setting makes it possible to explore whether income-targeting patterns extend beyond standard labor contexts, and to assess how such stopping rules relate to actual performance—something that cannot be observed directly in typical labor-supply environments.

In addition, the analysis contributes to the literature on skewness preferences. A large body of research shows that individuals value positively skewed outcomes in betting and investment settings (e.g., Golec and Tamarkin, 1998; Mitton and Vorkink, 2007; Barberis and Huang, 2008; Ebert, 2015; Dertwinkel-Kalt et al., 2024). These studies provide direct evidence on preferences for skewness in risky choice, whereas the present paper uncovers such preferences indirectly through stopping behavior. Building on the theoretical link recently formalized by Ebert and Köster (2024), we show empirically that stopping rules are systematically associated with the skewness of outcome distributions.

The empirical analysis reveals distinct patterns of reference-dependent stopping that systematically relate to both individual skill and the skill intensity of the environment. The patterns are more prevalent in environments where outcomes depend more heavily on skill, but less common among high-stakes players. Across all settings, higher-skilled players are more likely to exhibit gain-exit behavior, whereas lower-skilled players tend to display stop-loss patterns.

These behavioral tendencies are also associated with differences in session outcomes and profit distributions. Despite being more skilled and active overall, gain-exit players perform worse during periods of chasing, earning significantly lower expected profits within sessions when attempting to recover losses. Meanwhile, stop-loss players generate more right-skewed session-profit distributions, while gain-exit players display more left-skewed outcomes. Taken together, these patterns may suggest that highly skilled players pursue poker with implicit income targets, continuing play until reaching them. However, by doing so, they appear to undermine their short-term performance and unintentionally induce less favorable, left-skewed profit distributions.

The remainder of the paper is structured as follows. Section 2 describes the data. Section 3 outlines the empirical strategy, Section 4 presents the findings, and Section 5 discusses their broader implications and concluding remarks.

2 Data

Poker has evolved substantially from its early five-card draw origins, with variants such as Texas Hold'em now dominating both online and live play (Fiedler and Wilcke, 2011). This study uses tournament-level data purchased from the commercial provider HHSmithy, collected on the online platform PokerStars between November 2016 and February 2017. The analysis focuses on the *Heads-Up Sit-and-Go* (HUSNG) format of No Limit Texas Hold'em.

In a HUSNG tournament, two players compete in a self-contained match that begins as soon as both register at the same table. Each player pays an entry fee and receives an equal number of chips. Play continues until one player loses all their chips, at which point the other—who now holds all chips—wins the tournament and receives a prize equal to twice the entry fee.[1] Throughout the paper, we refer to each such tournament as a *match*, which constitutes the unit of observation in the analysis.

In Texas Hold'em, each player is dealt two private cards and can use up to five community cards—revealed sequentially and visible to both players—to form the

[1]PokerStars deducts a small fee, the rake, from the prize pool, which ranges from roughly 2% to 6% across game variants.

strongest possible five-card hand. In the No Limit version, players may bet any amount of their remaining chips at any time, including going “all-in,” creating a highly strategic environment characterized by continuous decision-making under uncertainty. The structure of Hold’em thus provides a rich context for studying dynamic behavior, as players must continually balance risk, potential reward, and their current position relative to the opponent.

Variations in starting chip endowments, decision times, and monetary stakes across match types influence the pace and intensity of play, potentially shaping players’ incentives to continue or stop. These contrasts provide a useful source of heterogeneity for studying behavioral patterns across environments. The analysis therefore draws on three datasets summarized in Table 1.

PokerStars offers two main match formats that differ in speed and structure, primarily through variation in the initial chip endowment. In *standard* (STD) matches, players start with 1,500 chips, while in *hyper turbo* (HT) matches they begin with only 500 chips. Both formats share the same blind structure,[2] but the smaller starting stacks in HT matches constrain strategic flexibility and generally lead to shorter matches. HT matches also impose tighter time limits for decision-making.[3] In addition to these speed variations, PokerStars offers matches at different stake levels. The present analysis includes both *micro-stake* (MS) $3.50 matches and *high-stake* (HS) $60 matches.

Table 1: Poker data included in this study

	Micro-stake	High-stake
Standard	Texas-STD-MS	
Hyper turbo	Texas-HT-MS	Texas-HT-HS

Previous empirical studies of poker behavior have typically relied on data from *cash games*, in which each hand is treated as a separate observation. While this approach provides large samples, it introduces several complications. In cash games, the optimal strategy in a given hand may depend on the history of previous play, particularly when opponents remain constant across hands. Moreover, players can leave or join tables at any time, and such entry and exit decisions may correlate with prior outcomes. For example, a player might quit after winning all the chips from a weak opponent who then leaves the table, but continue playing after losses if that weak opponent remains. These dynamics can induce systematic dependencies between past outcomes and continuation decisions.

By contrast, the Sit-and-Go format avoids such endogeneity. Players cannot leave before a match concludes, as doing so would mean forfeiting their remaining

[2]At the start of each match, the small blind is 10 chips and the big blind 20 chips, increasing at fixed intervals thereafter.

[3]STD matches allow 18 seconds per decision, whereas HT matches limit this to 12 seconds. If a player fails to act within the time limit, the system automatically checks (if possible) or folds.

chips. Each match ends deterministically when one player wins all chips, eliminating the possibility of mid-session cashing out. Importantly, when starting a new match against a different opponent, strategic considerations are reset, rendering each match independent of prior outcomes. This structure makes Heads-Up Sit-and-Go data particularly well suited for analyzing reference-dependent stopping behavior.

3 Methodology

This section introduces notation and definitions, and describes the methods to derive the main results of this study.

Sample selection and session definition

To identify stopping behavior reliably, it is crucial that players participate frequently enough to observe meaningful patterns in their continuation and stopping decisions. Moreover, players should make deliberate choices about when to stop, rather than being constrained by liquidity or external factors. For this reason, the analysis focuses on the most frequent players in the dataset, who tend to play multiple matches in close succession. Short intervals between matches make it unlikely that these players stop because they have exhausted their funds.[4] This group likely includes both professional players, who primarily play for profit, and individuals with problematic playing habits. As highlighted in previous research (Bjerg, 2010), these two characteristics are not necessarily mutually exclusive.

We define the sample as consisting of the percentile of players who participated most frequently during the observation period. This choice involves a trade-off: increasing the number of included players would introduce more noise, as less active participants may have played too few matches to infer their behavior reliably or may have been forced to stop due to limited funds; decreasing the number would reduce the sample size and the generality of the results. To verify robustness, we replicate the analysis using the top 200 most frequent players, and the main findings remain unchanged. Results for alternative thresholds are summarized in Appendix C.[5]

A key step in measuring reference-dependent stopping is defining when a player's reference point resets. The sequence of play between such resets is considered a single session. Following Eil and Lien (2014), we define sessions based on the time interval between consecutive matches. Specifically, consider player i who plays a match at time t_1. If the same player starts another match within one hour after

[4]Some players may still need to transfer additional funds between matches. The process of reloading an account on PokerStars is designed to be nearly instantaneous, though the need to do so may affect players' willingness to continue playing (Imas, 2016; Merkle et al., 2021). For further discussion, see Section 5.

[5]An earlier version of this paper also included the top 100 most frequent players, yielding consistent results.

completing the previous one, both matches are considered part of the same session, denoted s_1^i. If more than one hour elapses before the next match, session s_1^i is deemed to have ended, and the next match initiates a new session s_2^i. Hence, a session comprises consecutive matches separated by breaks of less than one hour, and a player ends a session when the current match is the last in that sequence.

For robustness, we also implement an alternative two-hour threshold and confirm that this adjustment does not substantially change the results (Appendix C). Table 2 reports summary statistics for sessions across datasets. Standard (STD) matches generally last about three times longer than hyper turbo (HT) matches, which partly accounts for the smaller number of matches per session observed in the STD format.

The approach proposed by Eil and Lien (2014) compared several bracketing intervals and found that shorter thresholds of one to two hours best reflect the time frame over which players integrate outcomes. Longer periods (such as six hours) produced more ambiguous classifications, as players may have already psychologically adjusted to prior results.[6] Adopting a one-hour benchmark, and testing robustness at two hours, therefore aligns with the evidence that players' reference points tend to reset over relatively short horizons.

Table 2: Statistics on matches and sessions - most frequent players

	Number of players	Total sessions	Mean sessions per player	Std. dev. sessions per player	Mean matches per session	Mean matches per player	Std. dev. matches per player
Texas-STD-MS	148	9,285	62.7	47.1	4.2	262.0	317.4
Texas-HT-MS	383	32,788	85.6	59.2	6.8	585.5	644.0
Texas-HT-HS	92	11,294	122.8	70.5	8.4	1025.4	719.1

Note: The table provides information on matches and sessions of the top percentile of most frequent players in each data set. Sessions are defined according to bracketing of one hour, i.e. breaks of at least one hour separate sessions from each other.

Identifying stopping behavior

To examine whether players' decisions to stop or continue depend systematically on their current session outcomes, we focus on two distinct strategy patterns: a *gain-exit strategy*, where players are more likely to stop when ahead, and a *stop-loss strategy*, where players are more likely to stop when behind. These patterns reflect reference-dependent behavior in which recent outcomes shape continuation decisions within a session.

[6]Indeed, we also observe some differences in performance and time until the next session when comparing session definitions based on one-hour versus two-hour breaks (see Section 4). These differences may reflect players' mental adjustments over time.

We classify stopping behavior by analyzing whether players decide to continue or end a session conditional on current profits. Specifically, we calculate cumulative session profits for each player and define a binary variable *behind*, which takes the value 1 when cumulative profits are negative and 0 otherwise.[7] This specification implicitly treats zero cumulative profit as the reference point. While reference points are likely heterogeneous across players (Kőszegi and Rabin, 2006; Buchanan, 2020), zero profits represent a natural and observable *break-even threshold*: they mark the transition between gains and losses and provide a salient, comparable benchmark across players.[8]

We then relate the variable *behind* to an indicator of whether the player ended the session. The variable *end_session* equals 1 if the current match concludes a session and 0 otherwise. For each player, we test whether stopping decisions depend significantly on being ahead or behind using a likelihood-ratio χ^2 test of independence between *end_session* and *behind*. This player-level approach accounts for individual heterogeneity and performs robustly with moderate sample sizes.[9] A player is classified as following a *stop-loss strategy* if the frequency of session endings is significantly higher when behind (i.e., players are more likely to stop after losses) and this relationship is significant at the 1% level.[10] Conversely, a player is classified as following a *gain-exit strategy* if session endings are significantly more frequent when ahead. Players for whom no significant relationship is detected are labeled as *unclassified*.

Finally, note that the analysis assumes players are aware of their current session outcomes. While platforms do not display cumulative session profits directly, players can observe their overall account balance at any time and readily compare it to the amount they had when starting to play. Entry fees are deducted at the beginning of each match, and winnings are credited immediately after, allowing players to track gains or losses relative to their starting balance with minimal cognitive effort. Session profits thus provide a salient and easily inferred reference for stopping decisions.

Variation of environments and player skill

The data provide a unique opportunity to examine how structural and individual factors relate to players' stopping behavior. Specifically, players in our sample face distinct environments differing in game speed and monetary stakes, and they

[7]Cumulative profits account for the platform's fees, so total profits within a session never sum exactly to zero.

[8]An alternative approach could consider each player's average session profit as their individual reference point. Classifications under this definition are highly correlated with the original method, and the results of this paper remain virtually identical.

[9]Simulations confirm that the likelihood-ratio χ^2 test performs comparably to a hazard model in this context and slightly improves classification accuracy for smaller samples. The main results remain unchanged when using a hazard specification.

[10]Robustness checks using a 5% significance threshold yield consistent results.

also vary systematically in skill. To account for these dimensions, we construct measures of the playing environment and player ability, which are later included as sources of heterogeneity and as controls in the analysis.

To distinguish between stake levels, we define a binary variable *high stakes*, which takes the value 1 for players participating in $60 matches and 0 for those in $3.50 matches. Differences in monetary stakes are substantial, as there are multiple intermediate levels between these two tiers. To account for differences in the relative importance of skill across poker formats, we further define a binary variable *higher skill dependence*, which equals 1 for standard (STD) matches and 0 for hyper turbo (HT) matches. The distinction reflects that outcomes in hyper turbo matches rely more heavily on chance due to shorter match lengths and smaller chip stacks. Using the best-fit Elo algorithm of Duersch et al. (2020), we quantify this difference and find that skill differences have a larger effect on outcomes in the STD format: a one-standard-deviation skill advantage yields a win rate of 53.6% in STD matches, but less than 51% in HT matches. Levene's tests confirm that the variance in winning probabilities is significantly greater in the STD datasets, supporting the interpretation that these environments are more skill-dependent.

Beyond characterizing environments, the Elo framework also provides an individual-level measure of player skill. The best-fit Elo model introduced by Duersch et al. (2020) is designed to adapt to different game formats and summarizes each player's expected win probability as a function of past performance against opponents of known strength, offering an intuitive and transparent proxy for skill.

As an alternative measure, we estimate player skill using a linear probability model with player fixed effects. In this specification, the binary outcome "won match" is regressed on player dummies and controls for game characteristics. The estimated player-specific fixed effects can be interpreted as skill parameters, since they capture average performance conditional on opponents' skill and contextual factors.

While both approaches produce highly correlated skill estimates and consistent results, we rely on the Elo ratings in the main analysis. The Elo framework is more conservative when estimating skill for players with relatively few observed matches, as it incorporates uncertainty directly into the rating update process. This is particularly important in our setting, where many players appear infrequently and thus provide limited data for reliable fixed-effect estimation. Because frequent players also face these less-active opponents, using Elo ratings ensures a more stable and comparable measure of skill across the full player network. The regression-based fixed effects are therefore used as robustness checks, confirming that the main findings are not sensitive to the specific skill measure.

Performance and continuation strategies

A central contribution of this paper is to examine whether the use of different continuation strategies—specifically, stop-loss and gain-exit behavior—relates to players' performance within sessions. The analysis tests whether being behind or ahead in cumulative session profits affects the probability of winning the next match, and whether this relationship differs between players classified by their stopping tendencies. Specifically, we analyze gain-exit players' performance when behind and stop-loss players' performance when ahead, as these conditions correspond to situations in which each group's continuation tendency is most relevant—paralleling the behavioral contrast between chasing and house-money dynamics.

To formalize this, we define the binary variable *behind_before_match*, which takes the value 1 if cumulative session profits up to the current match are negative and 0 otherwise. Analogously, *ahead_before_match* equals 1 if cumulative session profits up to the current match are positive and 0 otherwise.[11] The outcome variable won_j^i indicates whether player i won match j (1 = win, 0 = loss). We then estimate separate mixed-effects logistic regressions for stop-loss and gain-exit players, allowing for player-specific random effects to account for repeated observations:

$$\begin{aligned}\text{logit}(\Pr(won_j^i = 1)) = \beta_0 + \beta_1 stop\_loss^i + \beta_2 behind\_before\_match_j^i + \\ \beta_3(behind\_before\_match_j^i \times stop\_loss^i) + u^i,\end{aligned} \tag{1}$$

$$\begin{aligned}\text{logit}(\Pr(won_j^i = 1)) = \beta_0 + \beta_1 gain\_exit^i + \beta_2 ahead\_before\_match_j^i + \\ \beta_3(ahead\_before\_match_j^i \times gain\_exit^i) + u^i,\end{aligned} \tag{2}$$

where u^i denotes player-specific random effects. The models account for the binary nature of the dependent variable and the hierarchical structure of the data (multiple matches per player). Coefficients on the interaction terms capture whether being behind or ahead affects performance differently for stop-loss and gain-exit players.

We complement each model by including control variables for skill dependence and stake level (as defined in the previous subsection), as well as the current length of the ongoing session to account for potential fatigue. The models are estimated separately for the two strategy groups because the first match of each session, by definition, cannot be classified as ahead or behind, making joint estimation less meaningful. Robust standard errors are clustered at the player level.

Skewness of session profits

We assess how players' continuation strategies relate to the distribution of their session outcomes. For each player, we compute the skewness of session profits and evaluate how unusual this observed value is relative to simulated benchmark distributions. The simulations preserve each player's observed session lengths and individual win rates, but randomize the sequence of matches and outcomes 99 times. This procedure generates a reference distribution of session-profit skewness under random stopping, while maintaining players' empirical characteristics.

[11] Note that these two variables are perfectly negatively correlated for all matches except the first match of each session, where cumulative profits are zero and both variables take the value 0.

We then calculate the percentile rank of the observed skewness within this simulated distribution, which we refer to as the *skewness of session profits*. A value above the 50th percentile indicates more positive skewness than expected under random stopping, while a value below the median indicates more negative skewness. This measure provides a standardized comparison across players with different activity levels and serves as the dependent variable in later analyses linking individual skill to outcome asymmetries.

Relation between session outcomes and timing of the next session

In addition to continuation decisions within a session, we examine how outcomes at the session level relate to the timing of subsequent play. Specifically, we test whether the interval between sessions depends on whether players ended the previous session with a gain or a loss, and whether this relationship differs between stop-loss and gain-exit players.

We define the time-to-event variable as the number of hours between the last match of one session and the first match of the next. An event is recorded when a player begins a new session following a previous one. The binary variable $lost\_session_t^i$ equals 1 if player i ended session t with a negative cumulative profit and 0 otherwise. To model the timing of the next session, we estimate a Cox proportional hazards model of the form

$$\begin{aligned} h_i(t) = h_0(t)\exp(&\beta_1 lost\_session_{t-1}^i + \beta_2 stop\_loss^i + \beta_3 gain\_exit^i + \\ &\beta_4 lost\_session_{t-1}^i \times stop\_loss^i + \beta_5 lost\_session_{t-1}^i \times gain\_exit^i), \end{aligned} \tag{3}$$

where $h_i(t)$ denotes the hazard of starting a new session for player i at time t, and $h_0(t)$ is the baseline hazard. The coefficients on the interaction terms capture whether the association between prior outcomes and the timing of the next session differs across the two strategy types. A positive coefficient implies a shorter interval between sessions (i.e., faster re-entry).

We estimate the model both without and with additional control variables for skill dependence, stake level, and the session length of the preceding session. Robust standard errors are clustered at the player level.

4 Results

This section presents the empirical findings. We begin by classifying players according to the stopping strategies introduced in Section 3—the *stop-loss* strategy, where players are more likely to stop when behind, and the *gain-exit* strategy, where players are more likely to stop when ahead. We then examine how these behavioral patterns relate to game environments, individual skill, overall activity, and profits. Next, we analyze whether the strategies are associated with differences in within-session performance and the distribution of session outcomes. Finally, we assess how session results relate to the timing of players' subsequent sessions.

Classification of stopping strategies

Table 3 summarizes the distribution of player classifications across the three poker datasets. Most frequent players are classified as neither stop-loss nor gain-exit, forming the baseline group. The prevalence of each strategy varies moderately across formats: gain-exit behavior is somewhat more common in standard-speed environments, whereas stop-loss patterns are nearly absent at high stakes.

Table 3: Individual classification by poker version

	Stop-loss	Unclassified	Gain-exit
Texas-STD-MS	0.115	0.703	0.182
Texas-HT-MS	0.086	0.812	0.102
Texas-HT-HS	0.011	0.859	0.130

Note: Proportion of players in each behavioral category identified at the 1% significance level according to the classification procedure in Section 3.

To explore which characteristics predict these strategies, Table 4 reports logistic regressions comparing each type to the unclassified (baseline) group. Stop-loss behavior is significantly less prevalent in high-stakes matches and more prevalent in environments with higher skill dependence. Gain-exit behavior shows no systematic relationship with either stakes or skill dependence. Individual skill, however, exhibits a distinct pattern: less skilled players are more likely to follow stop-loss strategies, whereas higher-skilled players are more likely to exhibit gain-exit behavior. These findings are consistent across alternative thresholds and robustness checks (see Appendix C).

Table 4: Determinants of stopping strategies – Logistic regressions

	Stop-loss vs. Unclassified	Gain-exit vs. Unclassified
higher skill dependence	0.803**	0.214
	(0.314)	(0.369)
high stakes	-2.037**	0.143
	(1.021)	(0.360)
individual skill	-0.022**	0.018***
	(0.010)	(0.007)
constant	-2.135***	-2.218***
	(0.193)	(0.176)
Observations	545	572

Note: Logistic regressions with robust standard errors in parentheses. Dependent variable equals one if the player is classified as stop-loss (left column) or gain-exit (right column), respectively. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Result 1. *(a) Players with lower individual skill are more likely to follow stop-loss strategies.*
(b) Higher-skilled players are more likely to exhibit gain-exit behavior.

Overall activity and earnings

We next examine whether stopping strategies relate to overall engagement and profitability. Table 5 reports negative binomial regressions explaining the number of matches played. Gain-exit players participate in significantly more matches than unclassified players, whereas stop-loss players appear to play somewhat less frequently, although this difference is not statistically significant in the main specification. Results are broadly consistent across robustness checks.

Table 5: Frequency of play – Negative binomial regressions

	Stop-loss vs. Unclassified	Gain-exit vs. Unclassified
stop-loss	-0.107	
	(0.090)	
gain-exit		0.752***
		(0.148)
higher skill dependence	-1.233***	-1.100***
	(0.096)	(0.113)
high stakes	0.668***	0.594***
	(0.089)	(0.091)
individual skill	0.011***	0.008***
	(0.002)	(0.002)
constant	6.184***	6.197***
	(0.046)	(0.047)
Observations	545	572

Note: Negative binomial regressions with robust standard errors in parentheses. Dependent variable is the total number of matches played. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Result 2. *(a) Gain-exit players engage in matches significantly more frequently.*
(b) Stop-loss players tend to play less often, though the effect is not statistically significant.

Turning to profitability, Table 6 presents OLS regressions explaining total net profits (log-transformed). Neither stop-loss nor gain-exit players earn significantly different profits than the unclassified group. This pattern is consistent across all versions and robustness checks.

Result 3. *Neither stop-loss nor gain-exit players earn significantly different profits than unclassified players.*

Table 6: Net profits – Ordinary least squares regressions

	Stop-loss vs. Unclassified	Gain-exit vs. Unclassified
stop-loss	-0.007	
	(0.031)	
gain-exit		-0.028
		(0.023)
total matches	0.000	0.000
	(0.000)	(0.000)
higher skill dependence	-0.078*	-0.071*
	(0.041)	(0.039)
high stakes	-0.033	-0.001
	(0.135)	(0.111)
individual skill	0.004**	0.004**
	(0.002)	(0.002)
constant	8.627***	8.631***
	(0.029)	(0.021)
Observations	545	572

Note: Ordinary least squares (OLS) regressions with robust standard errors in parentheses. Dependent variable is the natural logarithm of total net profits after a constant shift so that the smallest profit equals one. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Performance within sessions

We next analyze whether these stopping strategies correspond to differences in performance within sessions, focusing on gain-exit players' performance when behind (consistent with chasing behavior) and on stop-loss players' performance when ahead (consistent with the house money effect). Tables 7 and 8 report odds ratios from the mixed-effects logistic regressions described in Section 3. Gain-exit players generally win more matches overall, but perform significantly worse when behind within a session. In contrast, stop-loss players show no systematic difference in performance when ahead compared to others. Accounting for the platform's entry fees, the estimated effect implies that expected profits for gain-exit players decline by roughly 30–50% when entering a match while behind. The result is robust across specifications but loses significance under two-hour session bracketing (Appendix C), suggesting that taking breaks of more than one but less than two hours may already allow players to reset mentally, mitigating the performance drop observed when continuing without such pause.

Result 4. *Gain-exit players perform significantly worse when behind within a session.*

Distribution of session profits

We now turn to the distribution of session outcomes. Following the simulation procedure described in Section 3, we compute each player's percentile rank of observed

Table 7: Performance regression – success when behind (gain-exit model)

	(1) without controls	(2) with controls
behind before match	1.005	1.007
	(0.008)	(0.008)
gain-exit	1.074***	1.052***
	(0.022)	(0.018)
behind before match × gain-exit	0.962**	0.963**
	(0.018)	(0.018)
higher skill dependence		1.138***
		(0.018)
high stakes		1.013
		(0.011)
match in session		1.000
		(0.000)
constant	1.080	1.063
	(0.007)	(0.008)
Observations	357,352	357,352

Note: Odds ratios from the mixed-effects logistic regression (1), estimated with and without additional controls. The first match of each session counts as not being behind. Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Performance regression – success when ahead (stop-loss model)

	(1) without controls	(2) with controls
ahead before match	1.000	1.003
	(0.007)	(0.008)
stop-loss	0.914***	0.907***
	(0.025)	(0.024)
ahead before match × stop-loss	0.988	0.989
	(0.030)	(0.030)
higher skill dependence		1.145***
		(0.018)
high stakes		1.005
		(0.011)
match in session		1.000*
		(0.000)
constant	1.097	1.078
	(0.007)	(0.007)
Observations	357,352	357,352

Note: Odds ratios from the mixed-effects logistic regression (2), estimated with and without additional controls. The first match of each session counts as not being ahead. Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

session-profit skewness relative to a simulated reference distribution under random stopping. Figure 1 shows a clear negative relationship between individual skill and the skewness of session profits: higher-skilled players exhibit more left-skewed outcome distributions. This association is statistically significant ($p < 0.001$) and replicates for alternative player subsets (see Appendix C). Consistent with the predictions of Ebert and Köster (2024), players following a stop-loss strategy display more right-skewed profit distributions than unclassified players, whereas gain-exit players show less right-skewed distributions than the other groups. These differences are confirmed by pairwise Mann–Whitney–U tests ($p < 0.001$).

Figure 1: Relation between individual skill and skewness of session profits (top percentile)

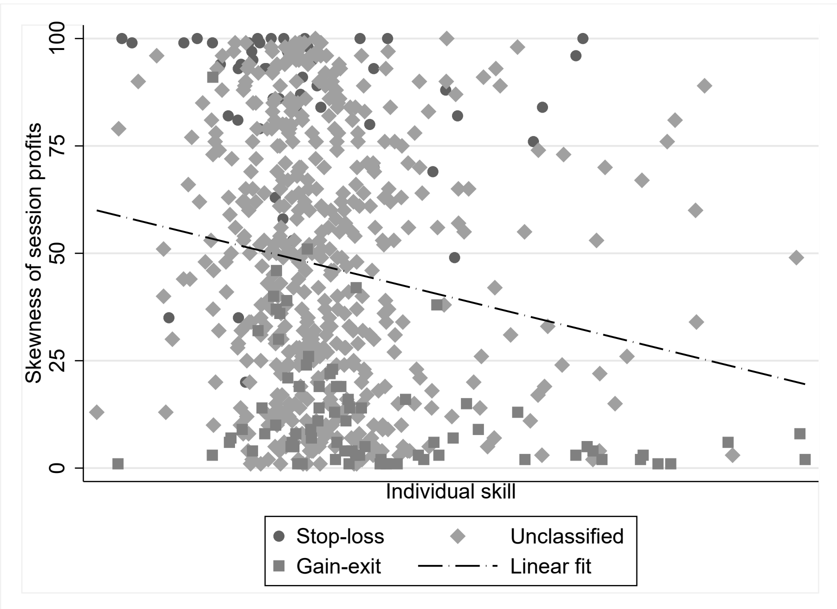

Note: The figure shows the relationship between individual skill and the skewness of session profits for the most frequent players in each poker version. Skewness is measured as the percentile rank of observed session-profit skewness relative to simulated distributions under random stopping. The negative relationship is statistically significant ($p < 0.001$).

Result 5. *(a) Higher-skilled players exhibit more left-skewed distributions of session profits.*
(b) Stop-loss players show more right-skewed profit distributions, whereas gain-exit players exhibit less right-skewed distributions relative to unclassified players.

Continuation across sessions

Finally, we examine whether session outcomes influence the timing of players' next sessions. Table 9 reports estimates from Cox proportional hazards models. Losing a session significantly delays the start of the next one on average. For stop-loss players, this delay is larger than for the unclassified group, while gain-exit players tend to resume play more quickly. In fact, the net effect of a lost session for gain-exit players is statistically indistinguishable from zero (Wald test, $p > 0.4$). These results remain consistent when including controls for skill dependence, stakes, and session length, and across alternative model specifications, but they are sensitive to the definition of session bracketing (Appendix C). Under a two-hour threshold, the earlier return of gain-exit players after losses is no longer statistically significant.[12]

Result 6. *(a) Gain-exit players' timing of return is largely unaffected by previous session outcomes.*
(b) Stop-loss players take shorter breaks after wins and longer breaks after losses.

Table 9: Timing of next session – Cox proportional hazards model

	(1) without controls	(2) with controls
lost session	-0.105***	-0.104***
	(0.012)	(0.012)
stop-loss	-0.045	-0.000
	(0.058)	(0.057)
gain-exit	0.017	0.059
	(0.042)	(0.044)
lost session × stop-loss	-0.079**	-0.089***
	(0.034)	(0.034)
lost session × gain-exit	0.097***	0.085***
	(0.029)	(0.030)
higher skill dependence		-0.158***
		(0.037)
high stakes		0.126***
		(0.037)
session length		-0.002**
		(0.001)
Observations	52,225	52,225

Note: Results from Cox proportional hazards models according to specification (3). Coefficients represent log hazard ratios, with positive values indicating shorter time until the next session and negative values indicating longer delays. Robust standard errors clustered at the player level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

[12] A likely explanation is mechanical: a substantial share of "quick returns" by gain-exit players occurs after breaks longer than one hour but shorter than two hours. When the one-hour rule counts these breaks as new sessions, gain-exit players appear to re-enter quickly; when the two-hour rule treats them as part of the same session, that apparent early return disappears.

5 Discussion

This paper provides new evidence on reference-dependent stopping in an environment where outcomes depend on both skill and chance. We find that stopping patterns are systematically related to individual skill and the skill intensity of the environment. Highly skilled players are more likely to stop when ahead (gain-exit behavior)—a pattern often associated with dynamic inconsistency in implementing stopping rules (Heimer et al., 2025)—whereas less skilled players are more likely to stop when behind (stop-loss behavior). Although gain-exit players typically win more often than they lose, their performance deteriorates during periods of chasing. This contrast indicates that reference-dependent stopping is a behavioral regularity that interacts with skill in complex ways and can, under certain conditions, impair performance.

A first interpretation draws on research on perceived control and overconfidence in skill-based domains. Gervais and Odean (2001) model learning under uncertainty about one's own ability, showing that successful individuals may become overconfident in environments combining skill and chance. Similarly, Bjerg (2010) argues that greater skill increases the cognitive complexity of evaluating control. In our context, skilled players may interpret losses as temporary deviations that can be reversed through superior play, reinforcing the tendency to continue after setbacks.

Furthermore, a dynamic mechanism may sustain such behavior once it emerges. In the Experience-Weighted Attraction (EWA) framework of Camerer and Hua Ho (1999), actions yielding higher payoffs become more attractive and are more likely to be repeated. Consistent with this idea, Campbell-Meiklejohn et al. (2008) find that unsuccessful chasers are less likely to chase in the future. For skilled players, frequent successful chases may therefore strengthen rather than discourage such behavior, consolidating the habit of continued play after losses.[13]

An additional mechanism that may underlie gain-exit behavior relates to motivated beliefs (Bénabou and Tirole, 2016). According to this theory, individuals derive utility not only from outcomes but also from maintaining self-enhancing views about their own ability. In this context, skilled players may prefer to preserve the belief that they are capable of playing profitably. Stopping after losses would conflict with that self-perception, while continued play allows them to reaffirm control and competence. Such belief maintenance may therefore reinforce gain-exit behavior.

Besides mechanisms that help explain why skilled players continue when behind, there are also potential explanations for why they may stop early when ahead. The pattern resonates with the literature on income targeting in dynamic labor supply. Camerer et al. (1997) first showed that New York City taxi drivers often stop once reaching a daily earnings goal, and subsequent work developed models of reference-

[13]Consistent with this interpretation, high-skilled players are substantially more likely to chase successfully. As shown in simulations (see Figure 5 in Appendix C), they are roughly three times more likely to chase successfully than an average player and about six times more likely than those in the bottom 5%.

dependent labor supply (e.g., Farber, 2008; Crawford and Meng, 2011; Thakral and Tô, 2021). In our data, gain-exit players—who are more active and typically winning players—appear to pursue similar implicit income targets, continuing play until they are ahead and then choosing to stop.

At the lower end of the profitability distribution, players are less likely to continue playing when behind, displaying stop-loss tendencies. A possible explanation relates to liquidity frictions: although these players are frequent participants and typically return to play after short intervals—making true constraints unlikely—their lower overall profitability compared to higher-skilled players may leave them with less money readily available on the platform. Continuing to play would then require transferring funds from external accounts—a step that transforms "paper" losses into realized ones. This mechanism resonates with the findings of Imas (2016) and Merkle et al. (2021), who show that realized losses increase risk aversion and reduce subsequent risk-taking.

The deterioration in performance during chasing episodes suggests that continuation after losses is accompanied by changes in play quality. Previous studies offer potential explanations for such behavioral shifts. Smith et al. (2009) and Suhonen and Saastamoinen (2018) document that individuals modify their strategies after losses, which could stem from shifts in risk attitudes or less disciplined decision-making following setbacks. A related experimental study by Grosshans et al. (2022) finds that trading decisions become less profitable after paper losses, driven by greater optimism and reduced belief sensitivity. A similar mechanism may operate here: poker players chasing their losses may attribute setbacks to bad luck and expect outcomes to revert in their favor (e.g., Tversky and Kahneman, 1971; Croson and Sundali, 2005), leading to more optimistic beliefs, riskier strategies, and ultimately lower expected profits.

Finally, the results contribute to research on how stopping rules shape the distribution of outcomes. A large body of work documents that individuals tend to prefer positively skewed payoffs in risky activities such as betting and investing (e.g., Golec and Tamarkin, 1998; Mitton and Vorkink, 2007; Barberis and Huang, 2008; Ebert, 2015; Dertwinkel-Kalt et al., 2024). In a related theoretical contribution, Ebert and Köster (2024) show that stopping behavior itself can systematically alter the skewness of returns: stop-loss behavior increases right-skewness, while gain-exit behavior induces more left-skewed outcomes. Consistent with this prediction, we find that players displaying stop-loss behavior have more positively skewed session profits, whereas gain-exit players exhibit more left-skewed outcomes. Given that higher-skilled players are more likely to adopt gain-exit patterns, they effectively generate less favorable payoff distributions.

Taken together, the findings raise the question of whether skilled individuals, who generally perform well, are more likely to adopt stopping patterns that inadvertently reduce short-term performance and generate less favorable profit distributions. More broadly, the results complement the findings of Heimer et al. (2025) on the welfare costs of dynamic inconsistency and excessive gain-exit behavior, highlighting how behavioral regularities can interact with skill to produce systematic and sometimes costly patterns of decision-making in dynamic environments. Future research could examine whether similar dynamics arise in other domains of repeated risk-taking—such as trading, entrepreneurship, or performance-based work—where outcomes are partly skill-driven yet shaped by feedback, beliefs, and self-perception.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the author used ChatGPT (OpenAI, GPT-5 model) to assist with language refinement, structural feedback, and improving the clarity and flow of the manuscript text. After using this tool, the author reviewed, edited, and verified all content, and takes full responsibility for the integrity and accuracy of the published article.

References

Barberis, N. and M. Huang (2008): "Stocks as lotteries: The implications of probability weighting for security prices," *American Economic Review*, 98, 2066–2100.

Barberis, N., M. Huang, and T. Santos (2001): "Prospect theory and asset prices," *The Quarterly Journal of Economics*, 116, 1–53.

Bénabou, R. and J. Tirole (2016): "Mindful economics: The production, consumption, and value of beliefs," *Journal of Economic Perspectives*, 30, 141–164.

Bjerg, O. (2010): "Problem gambling in poker: Money, rationality and control in a skill-based social game," *International Gambling Studies*, 10, 239–254.

Buchanan, J. A. (2020): "My reference point, not yours," *Journal of Economic Behavior & Organization*, 171, 297–311.

Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler (1997): "Labor supply of New York City cabdrivers: One day at a time," *The Quarterly Journal of Economics*, 112, 407–441.

Camerer, C. and T. Hua Ho (1999): "Experience-weighted attraction learning in normal form games," *Econometrica*, 67, 827–874.

Campbell-Meiklejohn, D. K., M. W. Woolrich, R. E. Passingham, and R. D. Rogers (2008): "Knowing when to stop: the brain mechanisms of chasing losses," *Biological Psychiatry*, 63, 293–300.

Coval, J. D. and T. Shumway (2005): "Do behavioral biases affect prices?" *The Journal of Finance*, 60, 1–34.

Crawford, V. P. and J. Meng (2011): "New York City cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income," *American Economic Review*, 101, 1912–32.

Croson, R. and J. Sundali (2005): "The gambler's fallacy and the hot hand: Empirical data from casinos," *Journal of risk and uncertainty*, 30, 195–209.

Dertwinkel-Kalt, M., J. Kasinger, and D. Schneider (2024): "Skewness preferences: evidence from online poker," *CESifo Working Paper*.

Duersch, P., M. Lambrecht, and J. Oechssler (2020): "Measuring skill and chance in games," *European Economic Review*, 127, 103472.

Ebert, S. (2015): "On skewed risks in economic models and experiments," *Journal of Economic Behavior & Organization*, 112, 85–97.

Ebert, S. and M. Köster (2024): "On taking a skewed risk more than once," *Available at SSRN 3731565*.

Eil, D. and J. W. Lien (2014): "Staying ahead and getting even: Risk attitudes of experienced poker players," *Games and Economic Behavior*, 87, 50–69.

Farber, H. S. (2008): "Reference-dependent preferences and labor supply: The case of New York City taxi drivers," *American Economic Review*, 98, 1069–82.

Feng, L. and M. S. Seasholes (2005): "Do investor sophistication and trading experience eliminate behavioral biases in financial markets?" *Review of Finance*, 9, 305–351.

Fiedler, I. and A.-C. Wilcke (2011): "The market for online poker," *Available at SSRN 1747646*.

Fraser-Mackenzie, P. A., T. Ma, M.-C. Sung, and J. E. Johnson (2019): "Let's Call it Quits: Break-Even Effects in the Decision to Stop Taking Risks," *Risk Analysis*, 39, 1560–1581.

Gervais, S. and T. Odean (2001): "Learning to be overconfident," *The Review of Financial Studies*, 14, 1–27.

Golec, J. and M. Tamarkin (1998): "Bettors love skewness, not risk, at the horse track," *Journal of Political Economy*, 106, 205–225.

Grosshans, D., F. Langnickel, and S. Zeisberger (2022): "Identifying Successful Trades: How Consistent Do Investors Act on Their Beliefs?" *Available at SSRN 2972112*.

Heimer, R., Z. Iliewa, A. Imas, and M. Weber (2025): "Dynamic inconsistency in risky choice: Evidence from the lab and field," *American Economic Review*, 115, 330–363.

Huang, Y. C. and S. H. Chan (2014): "The house money and break-even effects for different types of traders: Evidence from Taiwan futures markets," *Pacific-Basin Finance Journal*, 26, 1–13.

Imas, A. (2016): "The realization effect: Risk-taking after realized versus paper losses," *American Economic Review*, 106, 2086–2109.

Kőszegi, B. and M. Rabin (2006): "A model of reference-dependent preferences," *The Quarterly Journal of Economics*, 121, 1133–1165.

Lesieur, H. R. (1977): *The chase: Career of the compulsive gambler*, Anchor Press.

List, J. A. (2003): "Does market experience eliminate market anomalies?" *The Quarterly Journal of Economics*, 118, 41–71.

Locke, P. R. and S. C. Mann (2005): "Professional trader discipline and trade disposition," *Journal of Financial Economics*, 76, 401–444.

Merkle, C., J. Müller-Dethard, and M. Weber (2021): "Closing a mental account: The realization effect for gains and losses," *Experimental Economics*, 24, 303–329.

Mitton, T. and K. Vorkink (2007): "Equilibrium underdiversification and the preference for skewness," *The Review of Financial Studies*, 20, 1255–1288.

Siler, K. (2010): "Social and psychological challenges of poker," *Journal of Gambling Studies*, 26, 401–420.

Smith, G., M. Levere, and R. Kurtzman (2009): "Poker player behavior after big wins and big losses," *Management Science*, 55, 1547–1555.

Suhonen, N. and J. Saastamoinen (2018): "How do prior gains and losses affect subsequent risk taking? New evidence from individual-level horse race bets," *Management Science*, 64, 2797–2808.

Thakral, N. and L. T. Tô (2021): "Daily labor supply and adaptive reference points," *American Economic Review*, 111, 2417–2443.

Thaler, R. H. and E. J. Johnson (1990): "Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice," *Management Science*, 36, 643–660.

Tversky, A. and D. Kahneman (1971): "Belief in the law of small numbers," *Pediatrics*.

Zhang, W. and W. Semmler (2009): "Prospect theory for stock markets: Empirical evidence with time-series data," *Journal of Economic Behavior & Organization*, 72, 835–849.

A Measuring skill and chance

This appendix details how skill and chance are quantified across the different datasets. These measures provide the basis for comparing environments in terms of *higher skill dependence* and for deriving individual skill levels used in the main analysis. Table 10 summarizes the basic statistics of the datasets, restricted to *regulars*, defined as players who have participated in at least 25 matches.

Table 10: Statistics on matches, players, and regulars across poker datasets

	#Matches	#Players	#Regulars	Max Matches (Regulars)	Mean Matches (Regulars)	Median Matches (Regulars)
Texas-STD-MS	46,453	14,835	491	2,579	108.0	48
Texas-HT-MS	325,318	38,349	4,575	7,114	110.0	54
Texas-HT-HS	87,322	9,264	782	4,477	175.6	48

The best-fit Elo model

To measure skill heterogeneity and the role of chance, we employ the best-fit Elo model introduced by Duersch et al. (2020). The method estimates expected win probabilities based on past performance and calibrates a dataset-specific sensitivity parameter k, ensuring optimal predictive accuracy. For any match between players i and j at time t, the expected winning probability of player i is

$$E_{ij}^t = \frac{1}{1 + 10^{-(R_i^t - R_j^t)/400}},$$

where R_i^t denotes player i's current rating. After each match, ratings are updated according to

$$R_i^{t+1} = R_i^t + k \cdot (S_{ij}^t - E_{ij}^t),$$

with $S_{ij}^t = 1$ if player i wins and 0 otherwise. The optimal adjustment factor k^* minimizes prediction errors across all matches:

$$k^* = \arg\min_k \frac{1}{T} \sum_{t \in T} (S_{ij}^t - E_{ij}^t(k))^2.$$

Following Duersch et al. (2020), we restrict the analysis to players with at least 25 matches. Figure 2 shows that rating variability stabilizes beyond this threshold, suggesting that player ratings are well-calibrated from that point onward.

Interpreting Elo distributions

Once calibrated, Elo ratings summarize each player's relative skill within a dataset. The wider the rating distribution, the greater the predictability of outcomes, and thus the greater the role of skill relative to chance. Conversely, narrow distributions imply that outcomes depend more on luck. Figure 3 shows the unimodal distributions of Elo ratings for regular players, centered at zero for comparability.

Figure 2: Standard deviation of rating distributions by minimum matches per player

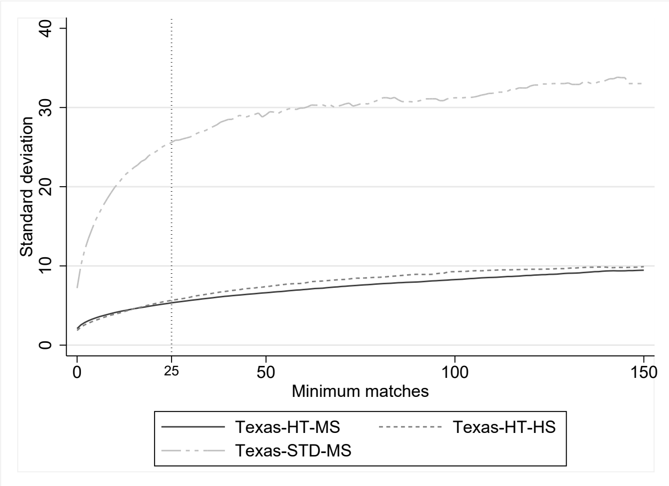

Note: The vertical dotted line marks the 25-match threshold used to define regular players.

Skill heterogeneity across environments

Table 11 summarizes the dispersion of Elo ratings across datasets. The standard deviation of ratings translates directly into differences in expected winning probabilities. A one–standard deviation skill advantage yields a 53.6% win probability in Texas-STD-MS, compared to only about 51% in the hyper-turbo formats. Levene's tests (Table 12) confirm that rating variances differ significantly between standard and hyper-turbo games.

Consistent with the main analysis, these results show that standard-speed matches are markedly more skill-dependent than hyper-turbo games. Differences between micro- and high-stake hyper-turbo formats are negligible, implying that higher stakes do not necessarily increase heterogeneity in play styles. This pattern is consistent with Siler (2010), who shows that high-stakes players tend to converge toward more optimal strategies, reducing variation and thus increasing the relative role of chance.

Complementary player fixed effects

For robustness, the main analysis complements Elo-based skill measures with player fixed effects estimated from a linear probability model that controls for opponent and game characteristics. Specifically, separate regressions are estimated for each poker version according to

$$won_{ij} = \alpha_i + \gamma_j + \varepsilon_{ij},$$

Figure 3: Rating distributions across poker formats

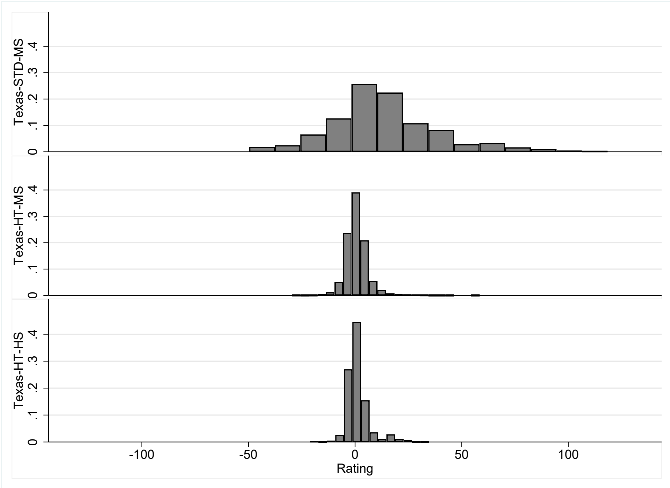

Note: Ratings are centered around zero by design. Only regulars (players with at least 25 matches) are included.

where won_{ij} is a binary indicator equal to one if player i wins match j, α_i represents player-specific fixed effects interpreted as individual skill, and γ_j captures game fixed effects. Standard errors are clustered at the player level.

The resulting player fixed effects are highly correlated with Elo ratings across all datasets (Pearson $r = 0.723$ in Texas-STD-MS, $r = 0.844$ in Texas-HT-MS, and $r = 0.811$ in Texas-HT-HS; all $p < 0.001$).

B Bias identification simulations

This section presents simulation results evaluating the classification procedure described in Section 3. The objective of this exercise is twofold. First, the simulations verify that the identification method correctly detects behavioral tendencies consistent with *stop-loss* and *gain-exit* patterns. Second, they test whether heterogeneity in players' underlying win rates could bias this classification, potentially generating mechanical associations between skill and stopping behavior. The results confirm that the procedure performs accurately and remains robust to moderate differences in success rates.

The identification strategy aims to distinguish individuals whose continuation decisions depend systematically on session outcomes. A *gain-exit* pattern corresponds to a higher probability of continuing play when behind—consistent with chasing losses—whereas a *stop-loss* pattern implies a higher probability of continuing when ahead, consistent with playing with house money. To assess whether the statistical tests used in Section 3 can recover these probabilistic patterns, we

Table 11: Summary statistics of Elo rating distributions

	Std Dev	Min	1%	99%	Max	p_{sd}	p_1^{99}	Rep_1^{99}	Rep_{sd}
Texas-STD-MS	25.3	-49.6	-44.5	91.9	107.4	53.6	68.7	3	89
Texas-HT-MS	5.3	-29.6	-11.0	18.9	54.6	50.8	54.3	61	1,777
Texas-HT-HS	5.6	-21.2	-9.1	23.2	32.5	50.8	54.6	55	1,777

Note: The table reports the standard deviation, minimum, 1st percentile, 99th percentile, and maximum of the rating distributions, which are centered at zero. Only regular players (those with at least 25 matches) are included. The table also translates ratings into directly interpretable quantities: p_{sd} denotes the winning probability of a player exactly one standard deviation better than their opponent; p_1^{99} denotes the winning probability of a 99th percentile player versus a 1st percentile player; Rep_1^{99} captures the number of matches required for the 99th percentile player to win more than half of the matches against the 1st percentile player with at least 75% probability; and Rep_{sd} reports the analogous number for a player who is exactly one standard deviation better than their opponent.

Table 12: Equality of rating variances (Levene's test)

	M0	M50	M10
Texas-STD-MS vs. Texas-HT-MS	0.000	0.000	0.000
Texas-STD-MS vs. Texas-HT-HS	0.000	0.000	0.000
Texas-HT-MS vs. Texas-HT-HS	0.677	0.319	0.342

Note: Reported values are p-values of Levene's test centered at the mean (M0), median (M50), and 10% trimmed mean (M10). Tests are applied to Elo rating distributions.

simulate data for artificial players who follow stochastic continuation rules rather than deterministic stopping points.

After each match, a simulated player's decision to continue is determined by a probability that depends on whether their cumulative profits are positive or negative relative to the start of the session. The probability of continuing when ahead is denoted p_a, and the probability of continuing when behind is denoted p_b. These parameters vary across simulated behavioral types:

$$(p_a, p_b) \in \{(0.85, 0.65), (0.8, 0.7), (0.75, 0.75), (0.7, 0.8), (0.65, 0.85)\}.$$

A player with $p_a > p_b$ exhibits a stop-loss tendency, whereas $p_a < p_b$ indicates a gain-exit tendency. When $p_a = p_b$, continuation decisions do not depend on prior outcomes. To examine the effect of individual success rates, players' match win probabilities vary across $p_w \in \{0.45, 0.525, 0.6\}$, corresponding approximately to the 5th, 50th, and 95th percentiles of observed win rates in the real data. Cumulative session profits, including platform fees, are updated after each match. Session continuation is determined probabilistically given p_a and p_b, with a maximum of 40 matches per session and 100 sessions per simulated player. Each simulation includes 10,000 players.

The generated data are then analyzed using the same identification procedure as applied to the empirical data. Specifically, we compute for each player the binary variable *behind*, indicating negative cumulative session profits, and *end_session*, which equals one if the player stops playing. Likelihood-ratio χ^2 tests are performed individually to evaluate the association between *end_session* and *behind*. Players who continue more often when behind ($p_a < p_b$) and whose relationship is significant at the 1% level are classified as displaying gain-exit tendencies; those who continue more when ahead ($p_a > p_b$) are classified as displaying stop-loss tendencies; and all others are treated as unclassified.

Tables 13 and 14 summarize the classification outcomes. A 10-percentage-point difference between p_a and p_b suffices to correctly identify roughly one-third of players at the 1% level and nearly two-thirds at the 5% level, with false classifications being extremely rare. When the difference increases to 20 percentage points, almost all players are correctly classified. As expected, when $p_a = p_b$, the misclassification rate aligns with the chosen significance threshold (about 1% and 5%, respectively).

Table 13: Simulations for stopping behavior identification at the 1% level

Winning %	Continuation % ahead	Continuation % behind	Stop-loss (%)	Unclassified (%)	Gain-exit (%)
45.0	85.0	65.0	96.55	3.45	0.00
45.0	80.0	70.0	35.48	64.52	0.00
45.0	75.0	75.0	0.58	98.85	0.57
45.0	70.0	80.0	0.00	64.89	35.11
45.0	65.0	85.0	0.00	1.95	98.05
52.5	85.0	65.0	98.32	1.68	0.00
52.5	80.0	70.0	38.77	61.23	0.01
52.5	75.0	75.0	0.49	98.94	0.57
52.5	70.0	80.0	0.00	59.41	40.59
52.5	65.0	85.0	0.00	1.09	98.91
60.0	85.0	65.0	98.76	1.24	0.00
60.0	80.0	70.0	39.41	60.59	0.00
60.0	75.0	75.0	0.47	99.00	0.53
60.0	70.0	80.0	0.00	61.30	38.70
60.0	65.0	85.0	0.00	1.77	98.23

Note: Simulations include 10,000 players, each completing 100 sessions of up to 40 matches.

Table 14: Simulations for stopping behavior identification at the 5% level

Winning %	Continuation % ahead	Continuation % behind	Stop-loss (%)	Unclassified (%)	Gain-exit (%)
45.0	85.0	65.0	99.14	0.86	0.00
45.0	80.0	70.0	58.74	41.26	0.00
45.0	75.0	75.0	2.52	94.88	2.60
45.0	70.0	80.0	0.00	40.26	59.74
45.0	65.0	85.0	0.00	0.29	99.71
52.5	85.0	65.0	99.65	0.35	0.00
52.5	80.0	70.0	62.88	37.11	0.01
52.5	75.0	75.0	2.33	94.96	2.71
52.5	70.0	80.0	0.00	34.99	65.01
52.5	65.0	85.0	0.00	0.16	99.84
60.0	85.0	65.0	99.80	0.20	0.00
60.0	80.0	70.0	63.34	36.66	0.00
60.0	75.0	75.0	2.32	95.23	2.45
60.0	70.0	80.0	0.00	37.03	62.97
60.0	65.0	85.0	0.00	0.34	99.66

Note: Simulations include 10,000 players, each completing 100 sessions of up to 40 matches.

C Additional analysis & Robustness

This section provides robustness checks for individual classification according to the identification strategy laid out in Section 3. In particular, we provide results for sessions being defined to allow for gaps of up to two hours between subsequent matches. We also include results for the top 200 most frequent players, as well as identification at the 5% level. Finally, for analysis that involves individual skill estimates, we complement Elo ratings from the main specification with estimates from fixed effect regressions for robustness (see Appendix A.) Table 15 summarizes statistics on sessions and matches for the different parameters. Table 16 provides the classifications for different parameter sets of the top percentile and the 200 most frequent players, respectively. Table 17 reports robustness checks for the logistic regressions estimating the effects of environments and individual skill on classification. Tables 18 and 19 provide robustness checks for results on frequency of play and net profits. Tables 20 and 21 present performance regressions for varying parameter settings. Figure 4 demonstrates the relation between individual skill and the skewness of session profits for the 200 most frequent players. Table 22 reports regressions on the time until the next session. Finally, Figure 5 depicts simulation results that corroborate the discussion of the main paper by providing evidence on the likelihood to chase unsuccessfully.

Table 15: Statistics on matches and sessions – most frequent players

	Number of players	Total sessions	Mean sessions per player	Std. dev. sessions per player	Mean matches per session	Mean matches per player	Std. dev. matches per player
Texas-STD-MS	200	10,739	53.7	43.8	3.9	211.4	285.9
Texas-HT-MS	200	20,768	103.8	68.0	8.1	845.2	808.1
Texas-HT-HS	200	15,230	76.2	66.4	7.3	557.5	652.6
Texas-STD-MS	148	7,153	48.3	32.3	6.9	262.0	317.4
Texas-HT-MS	383	25,169	65.7	42.6	11.1	585.5	644.0
Texas-HT-HS	92	7,931	86.2	46.1	13.5	1025.4	719.1

Note: The table provides information on sessions and matches under different parameters. The top section refers to the top 200 most frequent players, with sessions defined using one-hour bracketing. The bottom section refers to the top percentile of most frequent players, with sessions defined using two-hour bracketing.

Table 16: Individual classification by poker version

	Stop-loss	Unclassified	Gain-exit
Texas-STD-MS	0.110	0.745	0.145
Texas-HT-MS	0.065	0.785	0.150
Texas-HT-HS	0.090	0.845	0.065
Texas-STD-MS	0.176	0.601	0.223
Texas-HT-MS	0.141	0.697	0.162
Texas-HT-HS	0.033	0.728	0.239
Texas-STD-MS	0.101	0.750	0.149
Texas-HT-MS	0.081	0.843	0.076
Texas-HT-HS	0.000	0.891	0.109

Note: The table reports player classifications for the top percentile of the most frequent players, according to the identification strategy in Section 3. The top panel refers to the top 200 most frequent players, identification at 1%, and one-hour bracketing. The middle panel uses the 5% level with one-hour bracketing. The bottom panel applies the 1% level with two-hour bracketing.

Table 17: Individual classification robustness – Logistic regressions

	Stop-loss vs. Unclassified				Gain-exit vs. Unclassified			
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
higher skill dependence	0.731**	0.718**	0.747**	0.469	-0.490	0.047	-0.066	0.668**
	(0.362)	(0.288)	(0.337)	(0.330)	(0.347)	(0.320)	(0.489)	(0.276)
high stakes	0.204	-1.432**	-	-2.144**	-0.806**	0.306	0.240	0.227
	(0.379)	(0.607)	-	(1.012)	(0.352)	(0.288)	(0.396)	(0.354)
individual skill	-0.014*	-0.020**	-0.025***	-4.569**	0.023***	0.015**	0.027***	2.626*
	(0.008)	(0.008)	(0.008)	(2.185)	(0.006)	(0.006)	(0.008)	(1.508)
constant	-2.381***	-1.495***	-2.224***	-2.169***	-1.908***	-1.575***	-2.632***	-2.161***
	(0.299)	(0.156)	(0.196)	(0.186)	(0.212)	(0.150)	(0.201)	(0.179)
Observations	528	506	480	545	547	540	577	572

Note: Logistic regressions with robust standard errors in parentheses. The dependent variable is being classified as stop-loss (columns i–iv) or gain-exit (columns v–viii). $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 18: Frequency of play robustness – Negative binomial regressions

	Stop-loss vs. Unclassified				Gain-exit vs. Unclassified			
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
stop-loss	-0.306***	-0.138*	-0.149	-0.180*				
	(0.101)	(0.081)	(0.106)	(0.093)				
gain-exit					0.811***	0.538***	0.802***	0.756***
					(0.161)	(0.123)	(0.171)	(0.138)
higher skill dependence	-1.779***	-1.250***	-1.252***	-0.997***	-1.638***	-1.025***	-1.098***	-0.922***
	(0.090)	(0.100)	(0.094)	(0.089)	(0.105)	(0.130)	(0.108)	(0.095)
high stakes	-0.323***	0.634***	0.643***	0.709***	-0.242**	0.551***	0.581***	0.628***
	(0.099)	(0.095)	(0.089)	(0.091)	(0.101)	(0.094)	(0.092)	(0.093)
individual skill	0.014***	0.011***	0.010***	0.567	0.011***	0.008***	0.008***	0.821*
	(0.002)	(0.002)	(0.002)	(0.382)	(0.002)	(0.002)	(0.002)	(0.434)
constant	6.510***	6.200***	6.215***	6.253***	6.473***	6.203***	6.222***	6.245***
	(0.058)	(0.051)	(0.046)	(0.051)	(0.060)	(0.052)	(0.047)	(0.052)
Observations	528	506	562	545	547	540	577	572

Note: Negative binomial regressions on *total matches*, with robust standard errors in parentheses. Columns (i) and (v) refer to the top 200 most frequent players, identification at 1% and one-hour bracketing. Columns (ii) and (vi) refer to the top percentile, identification at 5% and one-hour bracketing. Columns (iii) and (vii) refer to the top percentile, identification at 1% and two-hour bracketing. Columns (iv) and (viii) refer to the top percentile, identification at 1% and one-hour bracketing, with skill from fixed effects. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 19: Net profits robustness – OLS regressions

		Stop-loss vs. Unclassified				Gain-exit vs. Unclassified		
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
stop-loss	-0.035	0.003	0.028	-0.010				
	(0.035)	(0.026)	(0.018)	(0.031)				
gain-exit					-0.037	-0.006	-0.037*	-0.019
					(0.025)	(0.023)	(0.022)	(0.022)
total matches	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
higher skill dependence	-0.030	-0.085*	-0.079*	0.011	-0.030	-0.078*	-0.078**	0.007
	(0.040)	(0.045)	(0.041)	(0.015)	(0.031)	(0.042)	(0.039)	(0.013)
high stakes	-0.005	-0.053	-0.035	-0.020	0.010	0.000	-0.016	0.016
	(0.052)	(0.150)	(0.131)	(0.130)	(0.052)	(0.114)	(0.111)	(0.104)
individual skill	0.004***	0.005**	0.005**	0.894**	0.004***	0.004**	0.004**	0.939**
	(0.002)	(0.002)	(0.002)	(0.433)	(0.001)	(0.002)	(0.002)	(0.439)
constant	8.612***	8.625***	8.617***	8.621***	8.613***	8.630***	8.628***	8.624***
	(0.028)	(0.033)	(0.029)	(0.033)	(0.021)	(0.022)	(0.020)	(0.024)
Observations	528	506	562	545	547	540	577	572

Note: OLS regressions on the natural logarithm of total net profits after a constant shift so that the smallest profit equals one; robust standard errors in parentheses. Columns (i) and (v) refer to the top 200 most frequent players, identification at 1% and one-hour bracketing. Columns (ii) and (vi) refer to the top percentile, identification at 5% and one-hour bracketing. Columns (iii) and (vii) refer to the top percentile, identification at 1% and two-hour bracketing. Columns (iv) and (viii) refer to the top percentile, identification at 1% and one-hour bracketing, with skill from fixed effects. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 20: Performance regression robustness – success when behind (gain-exit specification)

	(i)	(ii)	(iii)	(iv)	(v)	(vi)
behind before match	1.012	1.015*	1.008	1.010	0.003	0.005
	(0.008)	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)
gain-exit	1.084***	1.063***	1.053***	1.039***	0.089***	0.064***
	(0.024)	(0.020)	(0.017)	(0.015)	(0.023)	(0.019)
behind before match × gain-exit	0.958**	0.959**	0.962**	0.963**	0.963	0.965
	(0.019)	(0.018)	(0.016)	(0.015)	(0.023)	(0.022)
higher skill dependence		1.105***		1.141***		0.127***
		(0.018)		(0.018)		(0.016)
high stakes		0.978*		1.011		0.012
		(0.011)		(0.011)		(0.011)
match in session		1.000**		1.000		-0.000
		(0.000)		(0.000)		(0.000)
constant	1.077***	1.071***	1.079***	1.062***	1.080***	1.063***
	(0.07)	(0.009)	(0.007)	(0.008)	(0.007)	(0.008)
Observations	322,800	322,800	357,352	357,352	357,352	357,352

Note: Mixed-effects logistic regressions reporting odds ratios for specification (1); robust standard errors in parentheses. Columns (i) and (ii) refer to the top 200 most frequent players, identification at 1% and one-hour bracketing. Columns (iii) and (iv) refer to the top percentile, identification at 5% and one-hour bracketing. Columns (v) and (vi) refer to the top percentile, identification at 1% and two-hour bracketing.

Table 21: Performance regression robustness – success when ahead (stop-loss specification)

	(i)	(ii)	(iii)	(iv)	(v)	(vi)
ahead before match	0.997	1.000	0.998	1.001	0.999	1.002
	(0.008)	(0.008)	(0.008)	(0.008)	(0.007)	(0.008)
stop-loss	0.917***	0.905***	0.928***	0.924***	0.915***	0.912***
	(0.026)	(0.025)	(0.020)	(0.020)	(0.026)	(0.026)
ahead before match × stop-loss	1.004	1.006	1.014	1.016	1.009	1.012
	(0.031)	(0.031)	(0.025)	(0.025)	(0.032)	(0.032)
higher skill dependence		1.111***		1.143***		1.143***
		(0.018)		(0.018)		(0.018)
high stakes		0.973**		1.005		1.006
		(0.011)		(0.011)		(0.011)
match in session		1.000**		1.000*		1.000*
		(0.000)		(0.000)		(0.000)
constant	1.100***	1.093***	1.100***	1.081***	1.096***	1.077***
	(0.007)	(0.009)	(0.007)	(0.008)	(0.007)	(0.007)
Observations	322,800	322,800	357,352	357,352	357,352	357,352

Note: Mixed-effects logistic regressions reporting odds ratios for specification (2); robust standard errors in parentheses. Columns (i) and (ii) refer to the top 200 most frequent players, identification at 1% and one-hour bracketing. Columns (iii) and (iv) refer to the top percentile, identification at 5% and one-hour bracketing. Columns (v) and (vi) refer to the top percentile, identification at 1% and two-hour bracketing.

Figure 4: Relation between *individual skill* and skewness of session profits (200 most frequent players)

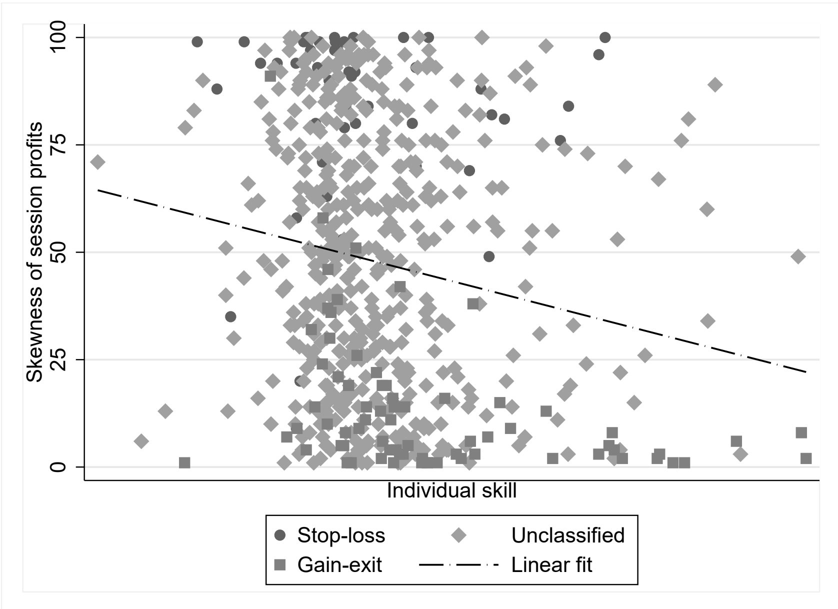

Note: The graph illustrates the relationship between individual skill and skewness of session profits for the 200 most frequent players in each poker version. To enable direct comparisons, players are represented using different symbols based on their strategy classifications. Skewness of session profits is determined through simulations in which session lengths and win rates remain fixed for each player, while match sequences and outcomes are randomly shuffled 99 times. We calculate the percentile rank of the observed skewness relative to the simulated reference distribution and refer to this measure as the skewness of session profits. The analysis reveals a statistically significant negative relationship ($p < 0.001$).

Table 22: Hazard model – time until next session (robustness)

	(i)	(ii)	(iii)	(iv)	(v)	(vi)
lost session	-0.107***	-0.110***	-0.103***	-0.102***	-0.096***	-0.096***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.014)	(0.014)
gain-exit	0.037	0.068	0.026	0.051	0.086*	0.129***
	(0.045)	(0.046)	(0.038)	(0.038)	(0.046)	(0.048)
stop-loss	-0.023	0.011	-0.022	0.013	-0.009	0.022
	(0.073)	(0.065)	(0.046)	(0.044)	(0.061)	(0.060)
lost session × gain-exit	0.116***	0.103***	0.093***	0.080***	0.037	0.018
	(0.029)	(0.029)	(0.026)	(0.026)	(0.039)	(0.039)
lost session × stop-loss	-0.143***	-0.151***	-0.093***	-0.099***	-0.074*	-0.076*
	(0.042)	(0.044)	(0.029)	(0.029)	(0.042)	(0.043)
higher skill dependence		-0.300***		-0.156***		-0.198***
		(0.039)		(0.037)		(0.033)
high stakes		-0.080**		0.118***		0.057
		(0.038)		(0.037)		(0.042)
session length		-0.002**		-0.002**		-0.000
		(0.001)		(0.001)		(0.001)
Observations	45,671	45,671	52,225	52,225	38,894	38,894

Note: Cox proportional hazards models per specification (3); coefficients are log hazard ratios (positive = shorter time to next session; negative = longer). Columns (i) and (ii) refer to the top 200 most frequent players, identification at 1% and one-hour bracketing. Columns (iii) and (iv) refer to the top percentile, identification at 5% and one-hour bracketing. Columns (v) and (vi) refer to the top percentile, identification at 1% and two-hour bracketing. Standard errors clustered at the player level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 5: Probability to chase unsuccessfully depending on maximum session length by individual success rate

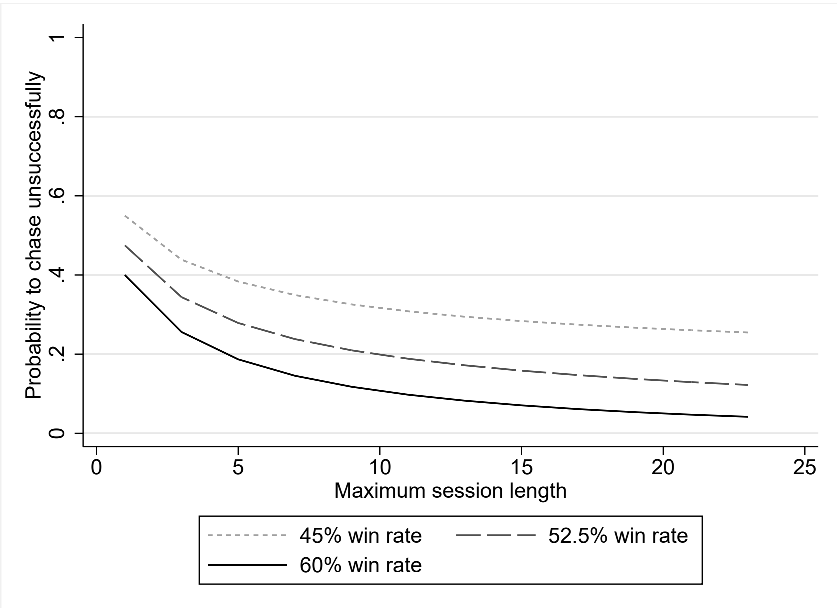

Note: The graph depicts the relationship between maximum session length and the probability to chase unsuccessfully depending on different levels of individual success rates, i.e., the probability to win a match $p_w \in \{0.45, 0.525, 0.6\}$. These correspond to the winning probabilities of the 5th, 50th, and 95th percentile players in the observed data. Results are based on simulations assuming that individuals play until they win one more match than they lost in the session (a successful chase) or reach the maximum session length without doing so (an unsuccessful chase).