COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 50 (2026) e70193

ISSN: 1551-6709 online
DOI: 10.1111/cogs.70193

A Metacognitive Appraisal of Quitting in Chess

Hariharan Purohit, Nisheeth Srivastava

Department of Cognitive Science, Indian Institute of Technology

Received 11 July 2025; received in revised form 27 January 2026; accepted 12 February 2026

Abstract

Decisions to stop or quit are often described as choices to switch to alternative activities once the current activity ceases to be rewarding. However, popular models of stopping—such as optimal foraging theory and recent metacognitive frameworks—fail to fully capture the nuanced differences in behavior between quitting and stopping. While quitting is closely related to stopping, it remains a phenomenologically distinct experience. The absence of a clear, separate definition for quitting motivates the present study. We investigate the contextual and noncontextual factors influencing quitting decisions among chess players, utilizing a large dataset of games from an online chess platform. Our analysis reveals that players tend to persevere in higher skill brackets and against stronger opponents when they are performing poorly. Additionally, a history of quitting increases the likelihood of quitting in future games, although recent quitting episodes can have protective effects. We also find that quitting influences subsequent behavior, with players often playing more games after a recent quit. We discuss these findings within the broader context of resource-rational and metacognitive approaches. Finally, we provide a metacognitive account of quitting decisions with the aim to derive better models of complex decision-making.

Keywords: Decision-making; Metacognition; Optimal stopping problem; Quitting decisions

1. Introduction

Stopping and quitting represent distinct classes of decisions that are applicable to any activity that a person is actively pursuing. In folk psychology, stopping and quitting are usually considered synonymous terms. However, these two types of decision can be distinguished by their phenomenological relevance. Stopping is characterized by the feeling of wanting to stop the activity and being able to stop the activity. Quitting, on the other hand, is characterized by wanting to continue on the activity and having to give up on the activity. In the latter

Correspondence should be sent to Nisheeth Srivastava, Department of Cognitive Science, Indian Institute of Technology, Kanpur, UP 208016, India. E-mail: nsrivast@iitk.ac.in

case, individuals tend to struggle with the decision and are visibly conflicted by their choices. This conflict can be characterized as wanting to continue exerting effort on the current task versus switching to another task. Currently, formal decision-making models do not explicitly consider this conflict when modeling stopping decisions nor do they differentiate between stopping and quitting.

This lack of distinction can be made clearer through a brief review of past literature on stopping and quitting. In information search paradigms, stopping decisions are defined as computing a stopping rule which can balance maximal rewards with minimal costs on the person's cognitive resources (Browne, Pitts, & Wetherbe, 2007; Branco, Sun, & Villas-Boas, 2012). An optimal solution to a stopping problem explicitly represents this balance (Charnov, 1976; Green, 1984). For example, in the secretary problem, the optimal strategy is to observe 37% of candidates and then select the next best, which maximizes the chance of choosing the top candidate (Ferguson, 1989).

Similarly, stopping rules have been linked to animal behavior. In optimal foraging theory, they are formulated to solve the problem of staying or giving up on a resource-rich patch. A typical stopping rule for an animal depends on the instantaneous rate of food consumption normalized by the search time for a particular patch (Charnov, 1976; Green, 1984). This simple rule accurately predicts that the animal "gives up" in a food patch when the consumption rate falls below an average capture rate, specific to that habitat and that animal.

However, these formalizations are not without limitations (Bugbee & Gonzalez, 2022). Since these models are developed in situ, the generated stopping thresholds are particular to that context, compromising their generalization. Recent studies have found that individuals dynamically adapt their thresholds based on changes in the task environment, especially when exploring their choice set (Bugbee & Gonzalez, 2024, 2025; Lee & Courey, 2021). While these models provide exemplars of how individuals generate decisions to stop, they limit their descriptions to utilitarian policies. Such descriptions imply that wanting to stop is built-in the decision itself. This leaves little room for variability in stopping across contexts.

Unlike stopping, quitting decisions have a normative expectation attached to them. This makes it cognitively effortful for the person to commit to the decision (Duke, 2022). If someone quits too early, they will miss out on the rewards that were "right around the corner." But if they persist too much in the face of constant adversity, they would incur opportunity costs (Alaoui & Fons-Rosen, 2021; Beshears & Milkman, 2011). Current theoretical frameworks have attempted to subsume these quitting behaviors under different conceptions.

The ability to quit unattainable goals, termed *goal disengagement*, has been associated with adaptive self-regulation of behavior. However, the ease of disengaging from a goal has been primarily attributed to factors outside the person's control, such as the reward structure of a task, social norms, or biological limitations (Wrosch, Scheier, Miller, Schulz, & Carver, 2003). Here, a quitting decision is equivalent to resolving a trade-off between persisting or completely disengaging from the activity. Statistically, this represents a modified explore–exploit trade-off (Sukhov, Dubey, Duke, & Griffiths, 2023), resolved by achieving an optimal balance between the expected future reward of quitting and the cost of lag in the decision.

Empirical data show that humans persist beyond the optimal quitting point (Alaoui & Fons-Rosen, 2021; Sukhov et al., 2023), in line with past research on optimal stopping rules. This operationalization of quitting shares certain commonalities with stopping, possibly indicating shared characteristics.

Another conception of a quitting decision is to define it as a solution to a strategy selection problem (Rieskamp & Otto, 2006). For example, quitting too early on an activity would represent selecting on a "satisficing" strategy (Ackerman, Douven, Elqayam, & Teodorescu, 2020; Artinger, Gigerenzer, & Jacobs, 2022). This strategy would mean that the person judges the current amount of information as a "good enough" stopping point. This stopping point is considered as near-optimal and represents a compromise due to limited cognitive resources. Such near-optimality in stopping thresholds has already been demonstrated in human observers who participated in artificial stopping problems (Bugbee & Gonzalez, 2025; Lee & Courey, 2021).

A recent modification of the strategy selection framework uses a metacognitive observer as an intermediary in the quit versus persist deliberation process (Ackerman et al., 2020). This observer balances the bounded cognitive resources with the cost of computing different strategies for a problem. The dominant proposal of a metacognitive observer involves both monitoring and control functions acting on the duration and cost of continuing to search for strategies (Nelson, 1990). Resource-rational meta-reasoning is a recent modification of this framework that proposes meta-reasoning as the process which resolves the strategy selection problem (Lieder & Griffiths, 2017). According to this theory, a resource-rational observer represents the computational cost of a strategy and its expected utility for the individual at a metacognitive level. The optimal strategy is then selected by deriving a meta-level value function across all strategies that are currently available to the individual. Within this approach, quitting can be framed as a resource rational strategy when the individual gives up on allocating further effort on the task, given low rewards.

1.1. Failure of current models of quitting

An ecologically valid example of quitting where existing models can be applied is loss-chasing behavior in gamblers. After accumulating a losing streak, gamblers choose to persist in a game by placing bets of increasing amounts. The expectation is that previous losses will be recovered in very few hands (Campbell-Meiklejohn, Woolrich, Passingham, & Rogers, 2008). This phenomenon, called "loss-chasing," is considered maladaptive due to gamblers' inability to consider the impact of previous losses while deciding on a strategy to pursue high-risk rewards.

Current models of stopping provide different predictions for loss-chasing. Broadly, both optimal stopping (Guan, Lee, & Vandekerckhove, 2015) and foraging (Green, 1984) models predict that a gambler decides to quit the table when their loss rate exceeds a predefined threshold. Specifically, foraging models predict stopping rules to be a function of both time and search cost. It is observed that when either one of these crosses a threshold, then the animal leaves the patch (Green, 1984). This would mean that gamblers will "cash out" when they run out of time or money. Studies have shown this is not always true, with gamblers persisting way beyond their stated thresholds for stopping (Browne, 1989; Campbell-Meiklejohn et al., 2008).

Recently, Bugbee and Gonzalez (2024, 2025) have shown that individuals' revealed preferences for stopping match poorly with predefined stopping thresholds, especially when they want to commit to them. Even when such thresholds are nonbinding, the stopping preferences deviate from optimal thresholds. This means that even when a stopping rule is predecided, people are unlikely to follow through on it. A plausible explanation for loss-chasing behavior could then be that it is a consequence of setting inconsistent stopping thresholds that do not affect their quitting decisions.

Utilitarian meta-reasoning would predict that gamblers choosing a loss-chasing strategy represent an optimal cutoff based on their available cognitive resources and sunk costs associated with quitting (Lieder & Griffiths, 2017). Quitting on this pursuit might instead represent a nonoptimal (but resource rational) strategy, with the resolution of an excruciating mental conflict as reward. Unlike previous models, meta-reasoning provides a flexible stopping rule depending on the gambler's history of quitting and persisting in a game. But these models are unable to account for the affective and behavioral consequences associated with loss-chasing, including how gamblers might form distorted beliefs about winning on the next bet, while facing a losing streak (Browne, 1989; Campbell-Meiklejohn et al., 2008).

Stopping and quitting are not clearly differentiated within these models, with most model parameters focused on utilitarian features of the problem space. Rarely do these models consider psychological variables (e.g., a person's knowledge about their stopping points) as model parameters. For example, the meta-reasoning model by Griffiths et al. (2019) involves computing a value, called "Value Of Computation (VOC)," that represents efficient usage of limited cognitive resources. Depending on the task, this value can also refer to the decision to allocate mental effort on it. In a similar sense for stopping tasks, the same optimal VOC could represent a successful strategy. Such a model parameter provides no theoretical distinction between stopping and quitting.

In this article, quitting refers to the revealed preference to forego allocating effort on a task even when the person wants to continue. This definition diverges from the utilitarian definitions mentioned above by using quitting behavior as a direct measure of the underlying process. Previous research has focused exclusively on laboratory-based stopping problems, which involve simpler decisions that occur on short timescales and are concerned with immediate rewards incurred on stopping at a particular point in the decision-making sequence (Guan, Lee, & Silva, 2014). However, quitting contexts involve complex decisions that occur on longer timescales and are concerned with consequences that are not easily accessible in laboratory experiments. Thus, a study of the properties of quitting is poorly matched to experimental methods. In this paper, a correlative analysis of a naturalistic domain is provided where individuals can intrinsically afford to quit in repeated episodes of the same activity, making their behavior amenable to empirical analysis (Kuperwajs & Ma, 2022). This behavior is assumed to be preceded by a push toward a decision to quit the effort allocation for the task. We provide an account that speculates that this decision involves latent metacognitive capabilities within the individual.

1.2. Quitting in chess

Chess is a popular domain for conducting ecologically valid studies of cognitive processes (Charness, 1977; De Groot, 1978). While prior work has examined stopping behavior in chess as a function of game variables (Chowdhary, Iacopini, & Battiston, 2023; Steyvers & Benjamin, 2019) or as a component of the planning process (Kuperwajs, Ho, & Ma, 2025; Russek, Acosta-Kane, van Opheusden, Mattar, & Griffiths, 2025), the present work directly investigates behavioral correlates of quitting and explores how these features are related to game factors.

Online chess offers numerous advantages as a domain for a careful study of quitting. Tracking performance of players across time can be done using different metrics. Player skill level is clearly shown through standardized Elo ratings that can be used to monitor a single player's performance over extended periods. The exact sequence of moves made by a player is also available, capturing the moment-to-moment fluctuations in choices that are made in response to the opponent. Furthermore, chess algorithms (such as *Stockfish*) can be used to quantify the board quality for each game, indicating how different strategies influence a player's decisions. The outcome of the game is not just represented by a win or loss, but also by how the game ends. Online chess databases provide various indicators of game termination, which, when combined with outcome indicators, can help classify games based on the player's decisions. One such indicator is game resignations, where a player decides to resign from the game rather than let it end in a win/loss. These advantages of online chess provide a robust foundation for our study of quitting behavior.

In chess, resignation decisions are initiated by the player but are not coded separately within game telemetry. Instead, resignation can be defined as the state of the game where the player decides to leave the game, provided they have not been forced into a state of checkmate. This decision translates into a series of steps taken by the player using the game interface to resign from the game. We define the decision to quit as equivalent to the outcome of resigning from the game. This differs from the decision to stop, which may not necessarily result in resignation. For example, a player can decide to stop the game by disconnecting from the server on which the game is being held. On the other hand, deciding to quit the effort of continuing in the current game would involve giving up on it. By restricting our analysis between resign and nonresign games, we can more naturally dissociate quitting behaviors from stopping behaviors, without compromising on the complexity of the phenomenon.

This study focuses on mapping the factors that affect the player's propensity to quit a game with respect to their overall preferences for quitting. Taking into account the normative expectations surrounding resignation in chess, it is likely that wanting to resign involves significant mental conflict for the player, extending well beyond any effect on their skill rating.

2. Method

2.1. Dataset

The raw dataset consists of online chess games sampled from a massively open-access online chess server *lichess.org*, with approximately 1 million games chosen from the year

2020. To capture a wide variance among players, this year was selected because of the substantial increase in the player base during the COVID-19 lockdown. The "Classical" category of games were chosen for the analysis because they provide a minimum move time of 30 min per player. The motive for choosing this category was to capture the temporally extended nature of quitting decisions, which cannot happen in game categories with durations lasting $\leq$ 10 min ("Blitz," "Rapid-Fire," etc.).

2.2. Preprocessing of data

To address the computational challenges associated with analyzing large-scale chess datasets, an extensive preprocessing pipeline was developed. Briefly, the pipeline included the extraction of games from PGN (Portable Game Notation) chess databases, which are downloaded from the *lichess* server. Inclusion criteria for the dataset consisted of games belonging to "Classical" category and those with "Normal" termination. The latter category refers to games having win or loss outcomes only. All other game categories and terminations (timeouts, disconnects, and draws) were excluded. The extracted games are then grouped by the player IDs and stored separately. A player's game roster includes games where they have played either as white or black. Multiple games played between two players are included in the roster for both players. Based on these steps, we arrive at a global sample of $N = 1{,}292{,}796$ games and $N = 14{,}300$ players. A random subset of players ($N = 250$) were chosen for further analysis. Two additional filters were applied at this stage: players with < 5 games and no resigns in their roster were removed.

The games for each player were analyzed using the *Stockfish v10.0* chess evaluation engine. Stockfish analyzes every half-move (*ply*) to calculate the quality of the board state, by specifying a numerical score (*centipawn*). This score quantifies the advantage or disadvantage of the half-move for the player. Besides the centipawn scores, the engine also provides scores that indicate distance to the next mate. Games that end in a checkmate contain mate scores instead of centipawn. To define a game as resignation, the game outcome should represent a loss for the player and the game should not have ended in a checkmate. A filter is applied to remove games that have mate evaluation scores that range from –5 to +5 instead of centipawn scores. The final sample size ($N = 218$) represents a reasonable compromise given the computational costs of analyzing such large data.

2.3. Measuring quitting

An online calculation of a player's propensity to quit a match can be defined as the rate at which the player resigns from their matches across the scale of a year. This measure is referred to as the player's Quit Rate, which is based on their current game history. In the context of a game, resignation is defined as a decision by the player to discontinue playing. Resignation can be considered as a within-game decision that contributes to the player's quitting propensity across games. This means that an increased frequency of resigns would reflect a greater propensity to quit. In the game interface, a resignation occurs when the player clicks the "Resign" button and then confirms this action during their turn. Fig. 1 illustrates a typical resignation scenario in which the player (playing as white) is at a greater disadvantage with

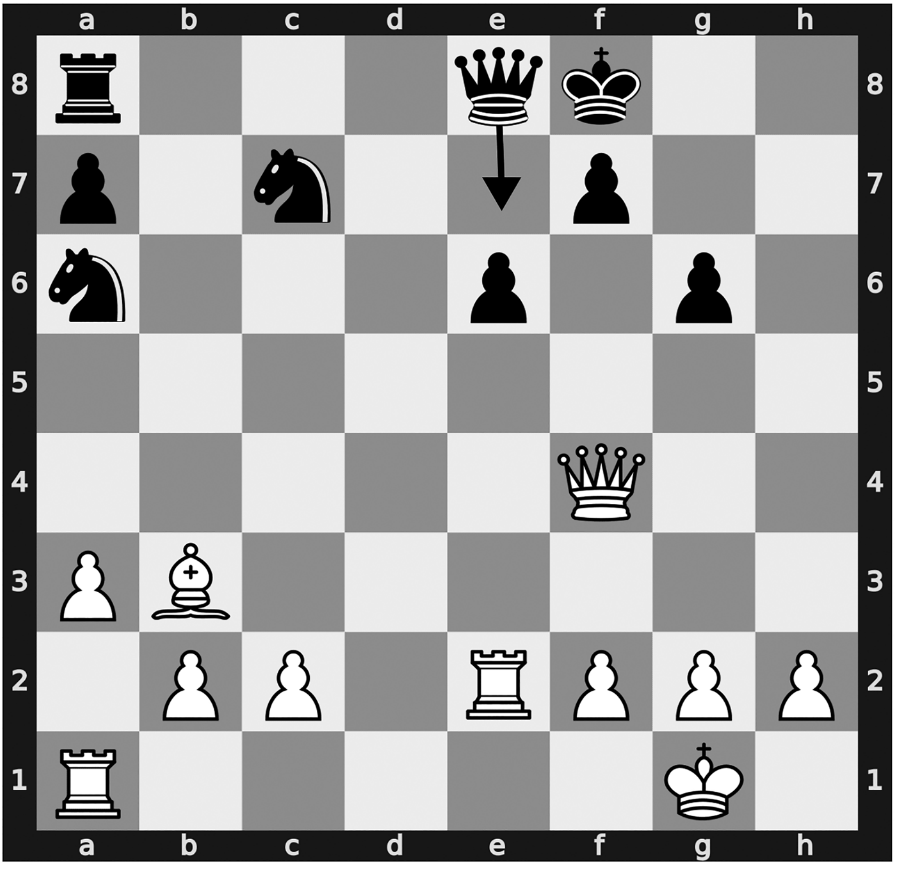

Fig. 1. An example of resignation scenario in chess.
Note. In this scenario, the white player resigns from the match. The black arrow represents the last move that was made by the opponent before the resignation.

respect to the board position, with no chance of recovery. While a resignation represents a within-game decision, the player could continue playing the game instead of resigning. This willingness to continue playing and yet, deciding to resign represents a core feature of quitting.

2.4. *Alternate specifications analysis*

In order to test the robustness of our results, we employed a recently developed technique, called *Specification Curve Analysis* (Simonsohn, Simmons, & Nelson, 2020). The primary goal of our specification analysis was to use an alternate specification of resignation in chess and evaluate three different analytical pathways in the processing pipeline. Across these three specifications, we defined resignation in chess differently from the original specification described above. To reiterate, resignation games represent those where the player

Table 1
Original and alternative reasonable specifications used to test the factors affecting quit rate and loss-hunting behavior

Decision	Original specification (Spec 1)	Alternative specifications (Specs 2–4)
Definition of resignation	Loss(Player) = 1 and Last Move(mate) = 0	Loss(Player) = 1 and Game Termination(Checkmate) = 1
Players with <= 5 games included?	No	No (Spec 2), Yes (Spec 3, 4)
Players with no resigns included?	No	No (Spec 2), Yes (Spec 3, 4)
Mate games included?	No	No (Spec 2, 3), Yes (Spec 4)

Note. Loss(Player) refers to when the player loses the game. Last Move(mate) refers to when the last move for the game was a distance to mate value. Game Termination(Checkmate) refers to when the game has terminated due to a checkmate.

loses the game and is not forced into a checkmate. Besides tracking the outcome indicators for a game, we can also track if the last move of the game represented a checkmate or not. This checkmate indicator evaluates the termination of the game due to a checkmate. Hence, resignation in a game can be defined as a loss outcome for a player and the game terminating without a checkmate, as classified by the checkmate indicator. The original specification excludes games that had moves evaluated as representing distance to a mate. Since this is not equivalent to a game terminating due to a checkmate, the alternate specification provides a more precise definition of player resignation. Player Quit Rates can then be calculated using this alternate specification.

Our specification analysis involves four descriptions (Table 1). Specification 1 (Spec 1) refers to the original dataset of 218 players and 406,098 games. Specification 2 (Spec 2) refers to a modified dataset with the same number of players as Spec 1 but having 386,760 games. The reduction in games with the same players could represent the alternative exclusion criteria (removing checkmate games) that provides a more specific definition of resigning. Specification 3 (Spec 3) refers to the same definition of resignation as Spec 2 but with the inclusion of players with < 5 games or players with zero resigns. Spec 3 contains a larger sample of data: $N = 8332$ players and $N = 615{,}168$ games. Spec 4 uses the same definition of resignation as the previous two and includes an additional subset of games that were removed from previous specifications. This subset includes noncheckmate games and games that had mate evaluations for the last move. This specification contains the entire dataset of players ($N = 14{,}300$) and games ($N = 1{,}969{,}312$) extracted from the chess server. The increase in number of games for Spec 4 compared to the raw dataset is due to the duplication of various games that involved two players repeatedly playing against each other. To ensure the calculation of player-specific quitting measures, such games were included as part of the roster of both players. Additionally, for Specifications 3 and 4, players with only a single game were excluded so that the measures could comply with our analyses.

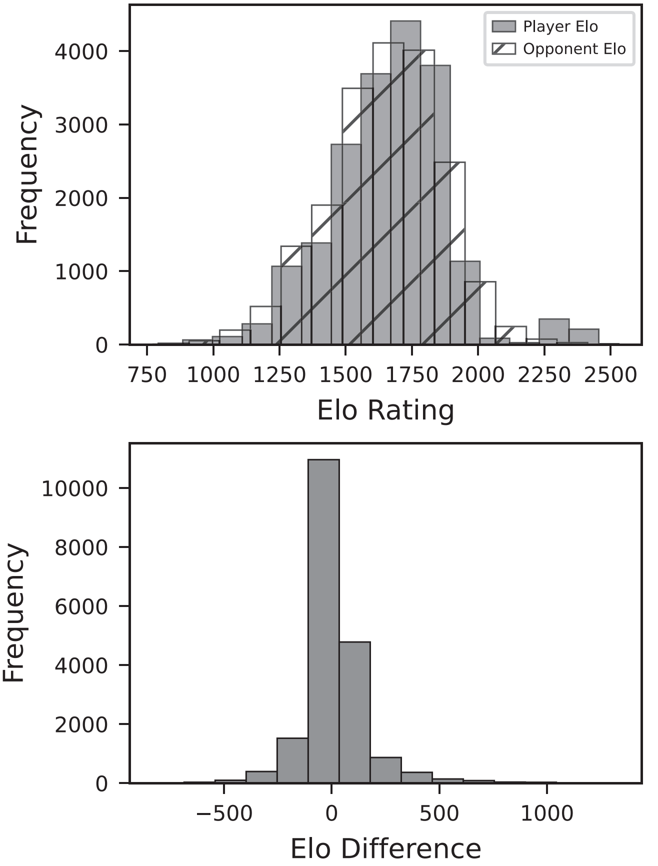

Fig. 2. Distribution of Elo ratings across players.
Note. Top panel: A histogram of Elo ratings for players and their opponents. Player ratings are shaded in light gray, while opponent ratings are shaded in stripes. Bottom panel: A histogram of differences in ratings between players and opponents.

3. Results

3.1. Dataset characteristics

The dataset demonstrates typical characteristics seen in chess, even with a smaller sample size. Fig. 2 reflects the typical properties matched to Elo that are present in online PvP (Player vs. Player) games, that is, the resulting differences in the skill of the players cluster around zero. A wide variety of game scenarios are represented in the dataset, as shown in Fig. 3. Fig. 4 visualizes the roster of a representative player with multiple resigns and nonresign matches, which is used to calculate their quit rate.

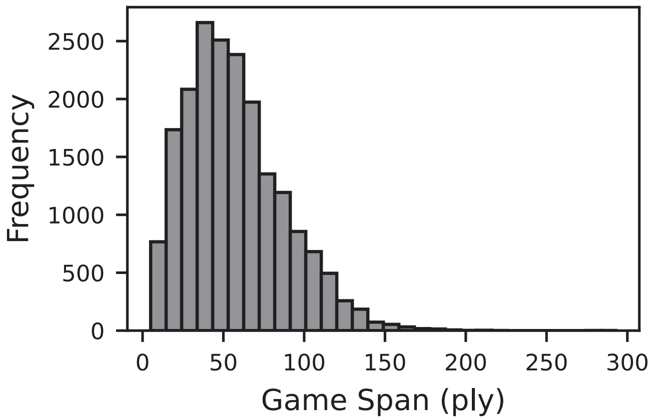

Fig. 3. Distribution of game length across players.
Note. Game span represents the total no. of ply moves made in the game, irrespective of player identity.

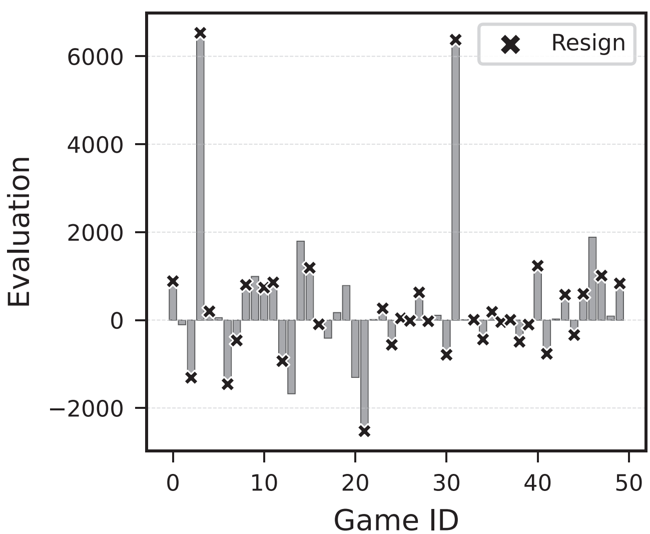

Fig. 4. Stockfish evaluations of games for a representative player.
Note. Stockfish evaluations are measured by centipawn units. Black crosses represent resigned games. Only the first 50 games for the player are shown here.

3.2. *Hazard model of quitting*

The aim of the paper was to empirically investigate how quitting decisions are influenced by both contextual game factors and past experiences of quitting in the game. This is accomplished by asking two questions: (1) How do game factors, such as skill differences between

Table 2
Mixed-effects Cox regression model with covariates for game factors and past quitting experiences

Covariate	$exp(\beta)$	95% CI	p-value
Evaluation	0.95	[0.93, 0.96]	<.001
Player ELO	0.87	[0.82, 0.92]	<.001
ELO difference	0.77	[0.75, 0.79]	<.001
Last Quit	0.93	[0.90, 0.96]	<.001
Total Quits	0.96	[0.93, 0.99]	.009
Evaluation × ELO Difference	0.97	[0.95, 0.99]	.003
Player ELO × ELO Difference	0.93	[0.91, 0.96]	<.001
Last Quit × Total Quits	0.98	[0.95, 1.01]	.3
ELO Difference × Last Quit	1.00	[0.97, 1.03]	.9
Evaluation × Last Quit × Total Quits	1.02	[1.00, 1.04]	.057

players and the quality of the board state, affect a player's decision to quit? and (2) Do players rely on their past experience of quitting in order to make quitting decisions?

To accomplish this, a typical game of chess can be framed as a game of survival for the player. Survival models specify the probability of surviving an event at a specific point in time (Kleinbaum & Klein, 1996), termed "survival rate." For a chess player, this event would be resigning from the match. Consequently, the probability that this event occurs at the next time marker t_{+1} given that the player has survived until t, can be defined. Hence, a hazard function can be calculated for the player resigning as a function of the resign tags and the final move number for each game. The outcome variable, termed the hazard rate, operationalizes a player's quit rate.

To capture the quit rate, a Cox proportional hazards regression model with mixed effects was formulated (Table 2). This model was used to assess the association between game factors, past quitting experiences, and quit rate, with a random effect for players to account for individual differences. The fixed effects included game factors such as the evaluation of the last move, the Elo rating of the player, and the difference in Elo ratings between the player and the opponent. Past quitting experiences were quantified by distance in time from the previous quit match to the current match (in seconds) and by the total number of quits. The estimated variance of the random intercept for the players was 0.325, indicating heterogeneity in the baseline hazard between the players. This suggests that individual differences had an effect on the variances of fixed effects.

3.2.1. *Impact of game factors on quitting*

In past research, human chess performance has been studied through the variability in opening sequences and win rates in different skill brackets (Chowdhary et al., 2023; De Marzo & Servedio, 2023), confirming the relation that win rate scales up with skill progression for a player. Similarly, we asked if there is a relation between skill progression and quit rate for players.

This can be answered with two game factors: evaluation of the last move in the game and the skill difference between a player and their opponent. The evaluation provided by

Stockfish represents the “ground truth” of the board state that a player approximates through their decisions (Holdaway & Vul, 2021). For quantifying player skill level, we also consider the skill difference (besides raw player ratings), as it represents a normalized metric for the player to judge their confidence of winning/losing in a match, w.r.t. the opponent.

Table 2 shows the results of a mixed-effects Cox regression model with player ID as random intercepts and adjustment for game factors and previous quitting experience. The results indicated that player skill level ($\exp(\beta) = 0.87$, 95 % CI: 0.82–0.92, $p < .001$), the difference in ELO ($\exp(\beta) = 0.77$, 95 % CI: 0.75–0.79, $p < .001$), and the evaluation of the last move ($\exp(\beta) = 0.95$, 95 % CI: 0.93–0.96, $p < .001$) predicted a 13 %, 23 %, and 5 % decrease in the risk of quitting a game, respectively. Highly skilled players tend to persist more in games, and playing against a significantly superior opponent strongly predisposes them to persist. A directionally similar but weaker effect is also observed when a player makes a significantly poor move. When difference in ELO and last move evaluation interact ($\exp(\beta) = 0.97$, 95 % CI: 0.95–0.99, $p = .003$), a 3 % decrease in the risk of quitting is predicted. Interestingly, after adjusting for player skill, the effect of skill difference persists ($\exp(\beta) = 0.93$, 95 % CI: 0.91–0.96, $p < .001$) through a 7% decrease in the risk of quitting. This interaction probably represents the inherent covariance between a player's ratings and the difference with regard to the opponent, directly factored into the Elo rating scheme.

3.2.2. *Quitting often makes quitting more likely*

Past experiences of quitting can be quantified in both time and frequency: *Last Quit* defines the time difference (in seconds) between the current game and the most recent quit game. *Total Quits* refers to the number of games the player has quit until the current game. Last Quit ($\exp(\beta) = 0.93$, 95 % CI: 0.90–0.96, $p < .001$) predicted a 7% decrease in the hazard rate of quitting, suggesting that the experience of having recently quit makes quitting in the current game less likely, keeping other factors constant. Similarly, Total Quits ($\exp(\beta) = 0.96$, 95 % CI: 0.93–0.99, $p = .009$) predicted a 4% decrease in the hazard rate of quitting. Interactions between past quitting experiences and game factors were not significant, suggesting that these factors independently affect quit rate.

3.3. *Loss-hunting due to quitting*

The complex nature of quitting decisions as a function of game factors and past quitting experiences should also translate into subsequent behavior. In order to recover the loss and embarrassment of having quit a match, players sometimes “tilt” toward playing more matches (Bonk & Tamminen, 2026). The influence of negative game performance on decision-making (Schüll, 2016) and on player emotions has been a persistent aspect of gameplay. This shift away from a player's normative game-playing behavior has been documented in poker players engaged in loss-hunting (Browne, 1989) and also in e-sports players (Wu, Lee, & Steinkuehler, 2021). We measured this behavior by calculating the number of matches played after the player had quit, relative to those where the player had not made this decision. We wanted to observe if chess players also express tilting behaviors as a hunt for more matches after racking up a substantial number of quits.

Table 3
Mixed-effects linear regression for deviation from base rate of play

Variable	Estimate	SE	95 % CI	p-value
Intercept	0.471	0.054	0.366, 0.576	<.0001
Player ELO	0.058	0.020	0.018, 0.097	.004
ELO difference	0.004	0.008	−0.012, 0.021	.594
Last Quit	−0.074	0.009	−0.092, −0.056	<.0001
Total Quits	0.056	0.011	0.034, 0.077	<.0001
Player ELO x ELO Difference	−0.017	0.006	−0.029, −0.005	.004
Total Quits x Last Quit	0.164	0.010	0.145, 0.182	<.0001

This loss-hunting behavior was measured as the deviation from a player's base rate of play after a resigned match. The base rate of play for a player was defined as the average number of games played over a week. The deviation from the base rate represents the total number of matches played over a week after a resignation, normalized by the base rate for that player. This deviation $D_{i,k}$ would indicate if the player i plays more (or fewer) matches after resigning relative to their average rate of play. The deviation for the k^{th} game provides a continuous rate-like metric to track the changes in the i^{th} player's gameplay.

Using this metric, a mixed-effects linear regression model was defined: $D_{i,k} \sim$ player Elo + Elo difference + Last Quit + Total Quits + $1|player$. The model follows a random-effects structure with base rate deviation as the dependent variable. Game factors (Player Elo and Elo difference) and past quitting experiences (Last Quit and Total Quits) are included as predictors in the model. Both predictors and the dependent variable were normalized by the z-score for the analysis. The player ID represents the random-effects variable in the model.

Table 3 shows the results of the mixed effects model. The Elo ratings of the players predicted the loss-hunting behavior of the players ($\beta = 0.058$, $SE = 0.020$, $p = .004$), indicating that the more skilled players show increased loss-hunting behavior. On the contrary, skill differences between player and opponent did not have an effect on loss-hunting behavior.

A player's previous quitting experience had a much greater impact on their loss-hunting behavior. Last Quit ($\beta = -0.074$, $SE = 0.009$, $p < .0001$) and Total Quits ($\beta = 0.056$, $SE = 0.011$, $p < .0001$) displayed main effects that indicate a case of loss-hunting behavior. When a player has quit very recently, their rate of play deviates to a greater extent from their base rate. Greater amounts of quits in the past matches also lead to greater amounts of loss-hunting. These effects indicate that a perpetual cycle of quitting can lead to increased volatility in the player's preferences for playing more games.

To understand how different factors contribute to the variance components of the model, the marginal and conditional R^2 values were calculated (Nakagawa & Schielzeth, 2013). The marginal R^2 was .20, indicating that fixed effects explained 20% variance in game play deviation after a resignation. Conditional R^2 was .52, indicating that 52% of the variance was explained when taking into account both fixed and random effects. The model remains robust despite the violations of the assumptions of homoscedasticity and normality (Schielzeth et al., 2020).

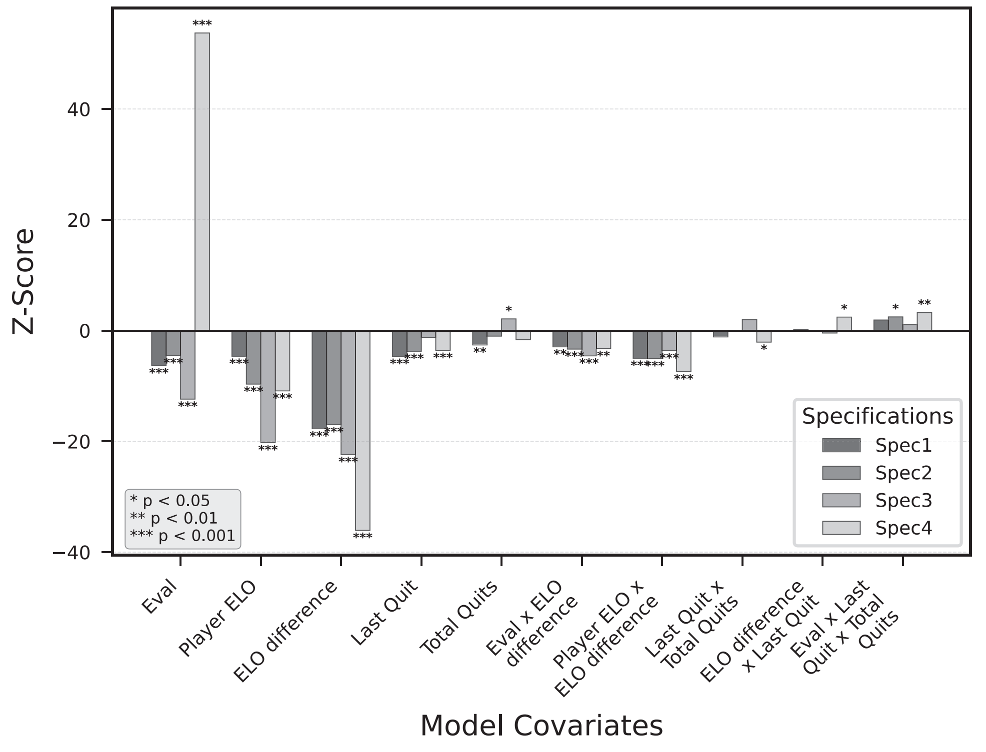

Fig. 5. Alternate specification plot for quit rate.
Note. Comparison of z-scores across four model specifications. Each specification is represented by a different shade of gray (Spec1: darkest; Spec4: lightest).
*$p < .05$, **$p < .01$, ***$p < .001$.

3.4. Alternative specification analysis results

Our alternative specifications were primarily done to check the robustness of our results across a greater number of players and games, while systematically reversing the decisions that were made to exclude data in our analyses. Across the specifications, results for quit rate and loss-hunting behavior displayed similar directional relations with game factors and past quitting experiences. Although specific differences do exist that reinforce our insights about quitting in chess players.

Beginning with the hazard model of quit rate (Fig. 5), larger effects are observed for player skill level and skill differences across specifications. Skill level shows a 22%, 27%, and 12% decrease in hazard of quitting for specifications 2, 3, and 4, respectively. Similarly, difference in Elo shows a 23%, 24%, and 27% decrease in the hazard of quitting for specifications 2, 3, and 4, respectively. These covariates provide robust estimates of game factors that affect the quit rate. Last move evaluation displays inconsistencies with a 4% and 10% decrease in quitting hazard for Specifications 2 and 3, while Specification 4 shows a 46% increase in hazard of quitting. The possible reason for this shift would be due to the high prevalence of mate games

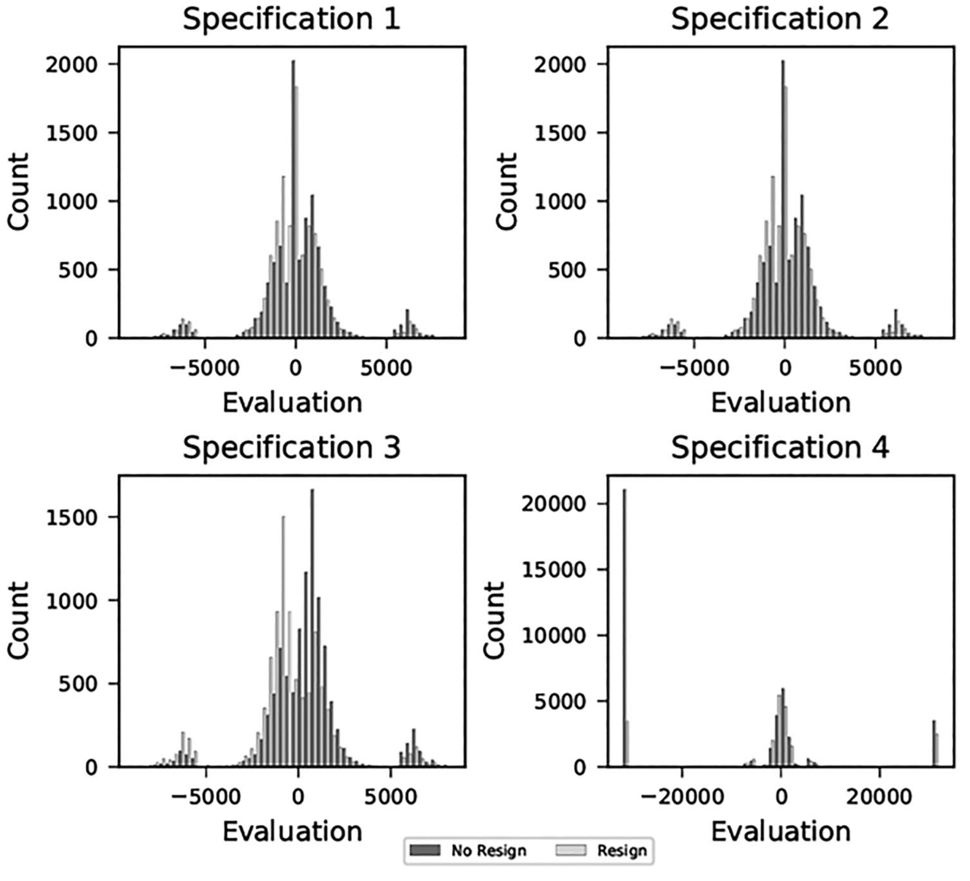

Fig. 6. Distribution of evaluation scores across the four alternative specifications.
Note. Each panel represents a different specification (Specifications 1–4). Within each panel, dark gray bars indicate games without resignation (No Resign) and light gray bars indicate games ending in resignation (Resign).

in Spec 4. Mate games were converted to centipawn games by assigning an extremely large number following this rule: $cp = \pm 32,000 - 2 \times$ mate value. The high prevalence of such values (Fig. 6), especially for noncheckmate games, could bias the results for interpretation.

Covariates associated with past history of quitting showed robust estimates: Last Quit showed a robust trend in decreasing the hazard of quitting. On the other hand, Total Quits showed weak estimates toward the hazard of quitting that did not reach significance in two specifications (Spec 2 and 4). Interactions between game factors and past quitting history were mostly robust, albeit weak predictors of the hazard of quitting.

Results for the linear mixed model of loss-hunting showed more robust relationships across the four specifications. Covariates for past quitting history (Last Quit and Total Quits) showed an increased trend across specifications indicating that a player's past experiences of quitting contributes toward loss-hunting behavior (Fig. 7). However, covariates based on game factors

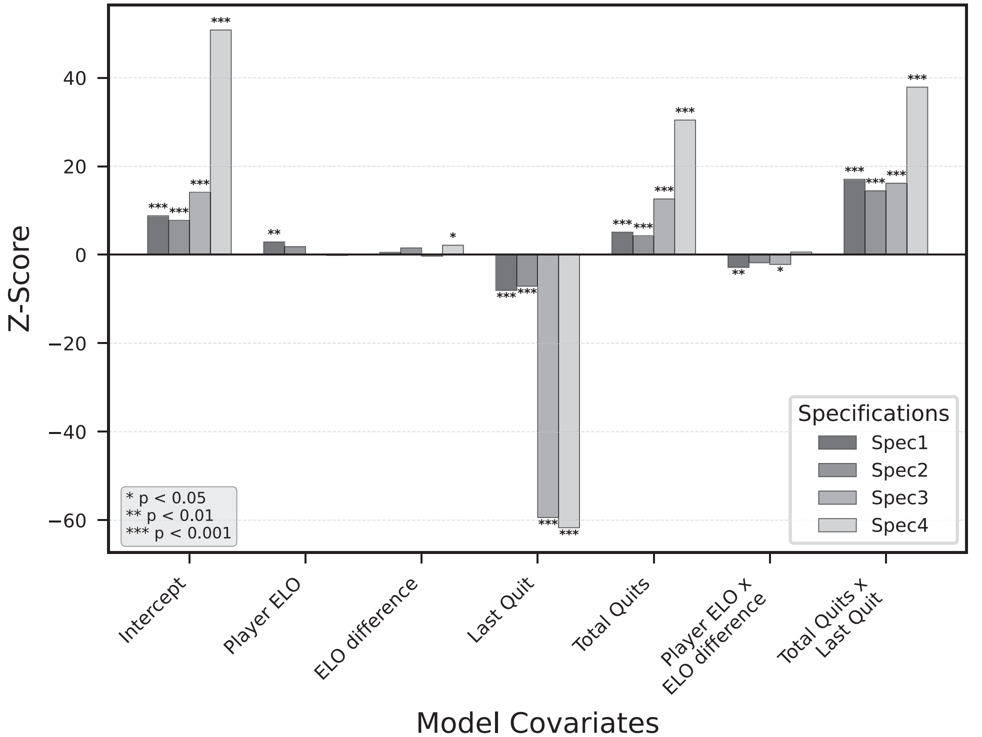

Fig. 7. Alternate specification chart for loss-hunting behavior.
Note. Comparison of z-scores across four model specifications. Each specification is represented by a different shade of gray (Spec1: darkest; Spec4: lightest).
*$p < .05$, **$p < .01$, ***$p < .001$.

(Player Elo and Elo difference) did not survive the alternative specification analyses, indicating that they may not contribute toward the loss-hunting behavior seen in chess players.

4. Discussion

The current paper was motivated by the phenomenological distinction between stopping and quitting that is felt in real-world contexts but has not been properly investigated in an empirical manner. Utilitarian models that have been proposed to explain these phenomena are unable to clearly differentiate between these types of phenomena. Due to the extensive work done on stopping decisions, this paper focused its attention on the lesser-known phenomena, quitting decisions. For this, chess is used as a middle ground between controlled laboratory experiments and real-world settings to investigate the factors behind quitting decisions and provide an aggregate measure of quitting derived out of the specific context of chess. We assume resignations in chess serve as a proxy for the decision to quit allocation of effort within the game. These within-game decisions drive a player's propensity to quit games in the long run. We defined this propensity to quit for a player as their Quit Rate. While Quit Rate provides an aggregate measure of quitting preferences across the player's game roster, it does not provide insight on how these decisions can impact future behavior. In order to check how the player behaves in future games, we also defined a deviation in the player's rate of play, after they quit from a game. We call this deviation, "loss-hunting" and evaluate the impact of past history of quitting on this behavior. Our results provide two novel insights into the quitting and gameplay preferences of a chess player. One, game factors such as player skill level, opponent strength, and the quality of the game state affect when players might decide to quit the game. These factors are found to drive the player toward persisting in adversarial game states. Additionally, a player's history of quitting also has a small, but significant impact on their quit rate. The influence of these factors reverses for a player's loss-hunting behavior, where a recent history of quitting seems to drive their hunt for more games. Player skill and opponent strength seem to not affect this behavior. We speculate about these results with reference to latent metacognitive capabilities that have been described in past research.

Chess involves sequential planning decisions (Kuperwajs et al., 2025) toward the goal of winning, which are generally measured by analyzing win rates (Chowdhary et al., 2023). The Quit Rate provides a similar player-centric measure of quitting decisions. This measure is affected by game factors such as the quality of the board state and the skill levels between players and their opponents. Players show increased persistence as their skill level increases. Also, players tend to persist longer against tougher opponents or when they are faced with a losing position. These results corroborate earlier work on risk preferences in chess players (Holdaway & Vul, 2021), which found that players prefer making riskier moves while in a losing position. They also found that highly skilled players tend to be more risk-averse and perform consistently regardless of the opponent's skill levels. These results might indicate that expert players show acute knowledge of the game board positions and the type of strategies in a game. One possible explanation of our results would be that these strategies might provide long-term benefits to the skilled player, such that persisting in the game at the current moment might benefit the player later in the game.

Quit Rate is also affected by the player's history of quitting. Having quit closer to their current match provides an adaptive shielding against future quitting and so does having a greater number of quits in past games. In order to explain these results, we can frame resignations as deliberations about effort allocation in the game. If the player decides to quit this deliberation, a resignation will be made and if they decide to delay it, they would persist in the game. Assuming this deliberation process involves considerable mental resources, delaying this decision by persisting in the task could represent a resource-rational strategy (Lieder & Griffiths, 2017; Lieder, Krueger, & Griffiths, 2017). One caveat is that no self-report data about chess players' stated preferences for persisting or resigning were collected that could corroborate our speculations. An alternative explanation for this behavior could be that players choose to persist in losing positions in order to learn new strategies from the opponent. Through this strategy, a player can turn the expected psychological loss incurred from past resigns into a learning opportunity. One way to test this possibility would be comparing the value of computation associated with time spent thinking in resign games and nonresign games. Recently, Russek et al. (2025) have demonstrated the overall benefit of computing multiple strategies, which is associated with an increased time spent thinking in online chess players.

Decisions to quit a match also have consequences for a player's future gameplay. As emotional guilt and embarrassment build up in a frequent quitter, the chances that they will engage in harmful tilting behavior in future games increase (Browne, 1989; Bonk & Tamminen, 2026). The cumulative failures can snowball into bigger lapses in decisions that could possibly lead to an even greater number of quits for a player. Our results indicate that players do fall into this perpetual cycle, where frequently quitting leads to a larger deviation in their base rate of play. We frame this increased engagement after a quit as the player hunting for games that potentially lead to greater losses. This loss-hunting behavior could possibly be driven by a satisficing strategy where the player intends to recover the mental losses by finding more games to play (Artinger et al., 2022; Bonk & Tamminen, 2026), instead of taking a break and resting. We also find that this behavior depends on a player's skill levels with skilled players engaging in greater loss-hunting behavior. The normative expectation of tilting (Bonk & Tamminen, 2025; Moreau, Sévigny, Giroux, & Chauchard, 2020) among players has been that as they become more skilled at playing the game, it becomes easier for them to recognize and prevent the negative consequences of tilting. But our data show evidence to the contrary, which might reflect players' need to "wipe off" the loss from the previous match, where they decided to quit. However, further losses that are incurred during loss-hunting could potentially lead to a perpetual cycle of tilt, which the player may have trouble getting out of (Bonk & Tamminen, 2026; Cregan, Toth, & Campbell, 2025).

The negative effect of a loss-hunting episode on a player's performance is also affected by their history of quitting. If a player has quit recently in their past matches, their loss-hunting behavior is greater and vice versa. In addition, the greater the number of quits that the player has, the greater their loss-hunting behavior. We speculate on two possible explanations for this behavior. One, players underestimate the impact of quitting experiences on their future gameplay behavior, choosing to keep on playing. Previous evidence with gamblers have shown that a small subset of players underestimate their subjective frequency of tilting episodes (Moreau et al., 2020), with minimal impact on their preferences for further gameplay. The other possibility is that players are acutely aware of their past quitting experiences and still choose to play more games. This is a likely possibility as online chess players do have access to their past game history, with their Elo rating acting as a summary statistic of their past quitting behavior.

We also checked if our results hold with a larger sample of players that display more variability in their distribution of resign and nonresign games. The alternate specification analyses considered a more specific definition of resignation in chess and found the results to remain robust across larger samples of games and players. Thus, the measures of quitting that we have proposed appear to account significantly for preference for quitting among chess players.

4.1. A possible role for metacognition in quitting

Quitting on the allocation of effort in a game of chess could involve metacognitive capabilities of a player. Below, we provide a speculative account for this based on the results of our study. The player's inability to consider the knowledge of their quitting experiences to decide on a strategy for future gameplay provides indirect evidence of metacognitive failure. This can be demonstrated through the results of loss-hunting behavior in players. The more recent is a quit in the player's history, more they seem to deviate from their weekly average rate of play. The phenomenological description of quitting indicates a gap between knowing when to withdraw and actually withdrawing from a task. This gap can appear as a lack of awareness for the players who frequently engage in loss-hunting. But this may not translate into their decision to quit the search for more games. Research on metacognition shows that metacognitive knowledge about future task performance can be dissociated from confidence reports on a task (Ackerman et al., 2020). This would mean that players may rely on their distorted beliefs about task performance over the explicit knowledge that is available to them about their recent history of gaming. Alternatively, if players are acting as rational meta-reasoners, their decision to seek out additional games may reflect a straightforward heuristic to continue playing due to the amount of mental effort already invested, as suggested by Lieder and Griffiths (2020). In other words, a player might choose to continue playing after resigning simply because they have already devoted significant time and energy on this chess server.

Another result that can result from the use of metacognitive capabilities is the effect of recent quitting experiences acting as a shield against further quitting. This could represent an example of meta-learning (Griffiths et al., 2019). In the meta-learning framework, an agent can use data from a limited set of tasks in order to learn and improve the hyperparameters of their learning algorithm. Assuming chess players act as rational meta-learners and applying this to our results, not quitting on games right after a recent quit could simply be a way to identify common mistakes that are made by the players. Leveraging these common mistakes by playing more games can prove advantageous for the player in the long run.

Our account of the role of metacognitive capabilities in quitting decisions is fraught with a few caveats. First, metacognitive knowledge about the task and confidence judgments about future performance require self-reports from players. We could not obtain these reports even after contacting players from our dataset on the lichess server. These self-reports could have provided data about the affective experiences of players on resigning and awareness about their loss-hunting behavior. Second, the persistence of players after a recent quit can have a simpler explanation that does not posit the use of any metacognitive capabilities. Players could simply continue playing under adversarial conditions due to the normative expectations associated with quitting a game. Quitting or resigning in competitive games is perceived as a disrespect to the opponent and lack of sportsmanship in the player. Players might just persist in order to avoid disrupting this normative expectation. Furthermore, the analysis of classical chess games may not extend to quitting scenarios seen in other types of game modes, where time pressures can significantly impact the utility of persisting or quitting. Lastly, the current correlational analysis cannot provide a mechanistic argument regarding a player's quitting experiences.

5. Conclusion

The current paper demonstrates a novel assessment of quitting behavior, separate from stopping behavior, in the popular naturalistic domain of chess. We investigate how this online measure of quitting impacts play behavior in chess players. Our correlational analysis speculates about the role of latent psychological factors that impact player behavior within the game over a long period of time. Our operationalization of quitting as a continuous measure embedded in player decisions goes beyond the post-hoc formalizations that are currently present in the related literature on stopping problems. Through this work, we appeal for a push toward more ecologically embedded studies of quitting behavior, with the hope of yielding decision-making models that possess greater phenomenological validity.

Conflict of interest

We have no known conflict of interest to disclose.

Data availability statement

The data and analysis scripts that can be used to reproduce the findings of this study are available in OSF at QuitChess project. The original chess data were derived from the following resource available in the public domain: database.lichess.org

References

Ackerman, R., Douven, I., Elqayam, S., & Teodorescu, K. (2020).Satisficing, meta-reasoning, and the rationality of further deliberation. In Shira Elqayam, Igor Douven, Jonathan St B. T. Evans, & Nicole Cruz (Eds.), *Logic and uncertainty in the human mind* (pp. 10–26). Routledge.

Alaoui, L., & Fons-Rosen, C. (2021). Know when to fold'em: The flip side of grit. *European Economic Review*, *136*, 103736.

Artinger, F. M., Gigerenzer, G., & Jacobs, P. (2022). Satisficing: Integrating two traditions. *Journal of Economic Literature*, *60*(2), 598–635.

Beshears, J., & Milkman, K. L. (2011). Do sell-side stock analysts exhibit escalation of commitment? *Journal of Economic Behavior & Organization*, *77*(3), 304–317.

Bonk, D., & Tamminen, K. A.. (2026). Exploring experiences of tilt and ragequitting in competitive and recreational videogamers. *Journal of Applied Sport Psychology*, *38*(1), 36–58, https://doi.org/10.1080/10413200.2025.2483688

Branco, F., Sun, M., & Villas-Boas, J. M. (2012). Optimal search for product information. *Management Science*, *58*(11), 2037–2056.

Browne, B. R. (1989). Going on tilt: Frequent poker players and control. *Journal of Gambling Behavior*, *5*(1), 3–21.

Browne, G. J., Pitts, M. G., & Wetherbe, J. C. (2007). Cognitive stopping rules for terminating information search in online tasks. *MIS Quarterly*, *31*(1), 89–104.

Bugbee, E. H., & Gonzalez, C. (2022).Deciding when to stop: Cognitive models of sequential decisions in optimal stopping tasks (Version 2). Carnegie Mellon University. https://doi.org/10.1184/R1/20492877.v2

Bugbee, E. H., & Gonzalez, C. (2024). Feedback promotes learning and knowledge of the distribution of values hinders exploration in an optimal stopping task. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 46.

Bugbee, E. H., & Gonzalez, C. (2025). Setting and adjusting thresholds in an optimal stopping task: Model predictions and empirical results. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 47.

Campbell-Meiklejohn, D. K., Woolrich, M. W., Passingham, R. E., & Rogers, R. D. (2008). Knowing when to stop: The brain mechanisms of chasing losses. *Biological Psychiatry*, *63*(3), 293–300.

Charness, N. (1977). Human chess skill. In Peter W. Frey (Ed.), *Chess skill in man and machine* (pp. 34–53).

Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, *9*(2), 129–136.

Chowdhary, S., Iacopini, I., & Battiston, F. (2023). Quantifying human performance in chess. *Scientific Reports*, *13*(1), 2113.

Cregan, S. C., Toth, A. J., & Campbell, M. J. (2025). The psychology of tilt: Understanding tilt and coping strategies in video games. *Cyberpsychology, Behavior, and Social Networking*, *28*(6), 408–416.

De Groot, A. D. (1978). *Thought and choice in chess*, volume *4*. Walter de Gruyter.

De Marzo, G., & Servedio, V. D. (2023). Quantifying the complexity and similarity of chess openings using online chess community data. *Scientific Reports*, *13*(1), 5327.

Duke, A. (2022). *Quit: The power of knowing when to walk away*. Penguin.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, *4*(3), 282–289.

Green, R. F. (1984). Stopping rules for optimal foragers. *American Naturalist*, *123*(1), 30–43.

Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M., & Lieder, F. (2019). Doing more with less: Meta-reasoning and meta-learning in humans and machines. *Current Opinion in Behavioral Sciences*, *29*, 24–30.

Guan, M., Lee, M., & Silva, A. (2014). Threshold models of human decision making on optimal stopping problems in different environments. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 36.

Guan, M., Lee, M. D., & Vandekerckhove, J. (2015). A hierarchical cognitive threshold model of human decision making on different length optimal stopping problems. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 824–829).

Holdaway, C., & Vul, E. (2021). Risk-taking in adversarial games: What can 1 billion online chess games tell us? In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 43.

Kleinbaum, D. G., & Klein, M. (1996). *Survival analysis: A self-learning text*. Springer.

Kuperwajs, I., Ho, M. K., & Ma, W. J. (2025). Heuristics for meta-planning from a normative model of information search. *Planning*, *1*(a2).

Kuperwajs, I., & Ma, W. J. (2022). A joint analysis of dropout and learning functions in human decision-making with massive online data. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 44.

Lee, M. D., & Courey, K. A. (2021). Modeling optimal stopping in changing environments: A case study in mate selection. *Computational Brain & Behavior*, *4*(1), 1–17.

Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, *124*(6), 762.

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1.

Lieder, F., Krueger, P. M., & Griffiths, T. (2017). An automatic method for discovering rational heuristics for risky choice. In *CogSci*.

Moreau, A., Sévigny, S., Giroux, I., & Chauchard, E. (2020). Ability to discriminate online poker tilt episodes: A new way to prevent excessive gambling? *Journal of Gambling Studies*, *36*(2), 699–711.

Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, *4*(2), 133–142.

Nelson, T. O. (1990). Metamemory: A theoretical framework and new findings. In Douglas Medin (Ed.), *Psychology of learning and motivation*, volume *26* (pp. 125–173). Elsevier.

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*(2), 207.

Russek, E. M., Acosta-Kane, D., van Opheusden, B., Mattar, M. G., & Griffiths, T. L. (2025). Time spent thinking in online chess reflects the value of computation. *Cognitive Science*, *49*(10), e70119.

Schielzeth, H., Dingemanse, N. J., Nakagawa, S., Westneat, D. F., Allegue, H., Teplitsky, C., Réale, D., Dochtermann, N. A., Garamszegi, L. Z., & Araya-Ajoy, Y. G. (2020). Robustness of linear mixed-effects models to violations of distributional assumptions. *Methods in Ecology and Evolution*, *11*(9), 1141–1152.

Schüll, N. D. (2016). Abiding chance: Online poker and the software of self-discipline. *Public Culture*, *28*(3), 563–592.

Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2020). Specification curve analysis. *Nature Human Behaviour*, *4*(11), 1208–1214.

Steyvers, M., & Benjamin, A. S. (2019). The joint contribution of participation and performance to learning functions: Exploring the effects of age in large-scale data sets. *Behavior Research Methods*, *51*, 1531–1543.

Sukhov, N., Dubey, R., Duke, A., & Griffiths, T. (2023). When to keep trying and when to let go: Benchmarking optimal quitting. *PsyArXiv*.

Wrosch, C., Scheier, M. F., Miller, G. E., Schulz, R., & Carver, C. S. (2003). Adaptive self-regulation of unattainable goals: Goal disengagement, goal reengagement, and subjective well-being. *Personality and Social Psychology Bulletin*, *29*(12), 1494–1508.

Wu, M., Lee, J. S., & Steinkuehler, C. (2021). Understanding tilt in esports: A study on young league of legends players. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1–9). Yokohama, Japan: ACM.