A metacognitive appraisal of quitting

Hariharan Purohit (hariharan22@iitk.ac.in)
Department of Cognitive Science, Indian Institute of Technology, Kanpur
India

Nisheeth Srivastava (nsrivast@iitk.ac.in)
Departments of Computer Science and Cognitive Science, Indian Institute of Technology, Kanpur
India

Abstract

Stopping decisions are frequently modeled as decisions to switch to alternative activities once the current activity stops being adequately rewarding, such as in optimal foraging theory, as well as more recent metacognitive models. However, the sense of stopping and making decisions in such frameworks is highly platonic, with both decisions and stopping actions occurring instantaneously. In contrast, the phenomenology of quitting actions that one is undertaking appears to be temporally extended and metacognitively challenging. We study the metacognitive covariates of quitting decisions made by chess players using a large database of chess games sourced from an online chess portal. Our analysis reveals that players tend to persevere when they are playing against stronger opponents and after having played poor moves. We also find that a history of quitting games makes players more likely to quit in future games, but that having recently quit in a game offers some protective effect against quitting. Finally, we find that quitting a game makes it more likely that a player will play a game again soon. We place these results in the context of modeling quitting as a metacognitive choice affected by multiple competing goals.

Keywords: decision-making; metacognition; optimal stopping problem, quitting decisions; resilience

Introduction

Quitting an activity that one is actively pursuing represents a phenomenologically distinct class of stopping decisions. In formal decision-making models, quitting is synonymous to stopping decisions, that are usually framed as deciding on a stopping rule for information search that maximizes reward and minimizes cost (G. J. Browne, Pitts, & Wetherbe, 2007; Branco, Sun, & Villas-Boas, 2012). It has long been recognized that an optimal solution to this optimal stopping problem (OSP) should strike a balance between these two (Charnov, 1976; Green, 1984).

In optimal foraging theory, the *fons et origo* of the formal study of stopping decisions, stopping decisions are framed as a problem of persisting or giving up in a food-rich patch for an animal. The decision to 'give up' on a patch is dependent on the rate of food consumption and the time spent searching within the patch (Charnov, 1976). Optimal foraging predicts that the animal should 'give up' on a patch once the rate of capturing food in a patch drops down to the average food capture rate within the area.

However, optimal foraging and its successor theoretical frameworks remove the complexities of real-life decision making by not incorporating the role of experience that can lead to changes in decision thresholds. They also have low generalizability to contexts beyond the specific task which are used to generate static thresholds (Bugbee & Gonzalez, 2022).

A real-life example where these frameworks fail to predict the optimal stopping rule is loss chasing in gamblers. Gamblers tend to recover their losses by placing bets of increasing sizes with the expectation that they will probably win this hand. Loss chasing is considered maladaptive because gamblers tend to underestimate their previous loss experience when deciding on a strategy to pursue high-risk rewards (Campbell-Meiklejohn, Woolrich, Passingham, & Rogers, 2008).

Optimal stopping and foraging theories do not account for the loss-chasing nature of gambling behavior. Both these models would predict that given a consistent losing streak, a gambler should set their stopping thresholds at a boundary where their reward rate would fall below expected future reward (Guan, Lee, & Vandekerckhove, 2015). Since gamblers do loss-chase, both these models would fall short of explaining this behavior unless additional parameters are introduced.

In other domains of life, individuals struggle with the problem of persevering in unfavorable tasks compared to favorable tasks. Research on strategy selection (Lieder & Griffiths, 2017; Lieder, Krueger, & Griffiths, 2017) has provided a different solution from classical stopping models. Inspired from bounded rationality, this research argues that humans select strategies that balance the cost of computation with the current availability of cognitive resources. This selection process is termed as meta-reasoning, that is, assessing how good or bad a particular strategy is. Meta-reasoning about when to stop would require deliberating over the optimal stopping strategy that balances the deliberation time with existing resources.

Returning to the example of loss-chasing gamblers, utilitarian meta-reasoning would predict that gamblers choose to continue gambling after a string of losses as this strategy still allows them to recover some portion of their losses ('high-risk, high-reward') given the reduced cognitive resources under which they would be operating. Qualitative evidence of loss chasing has shown how gamblers mention the hope of 'hitting the jackpot' on their next bet, even in the face of imminent doom of financial liabilities associated with their losses (B. R. Browne, 1989; Campbell-Meiklejohn et al.,

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

2008).

The meta-reasoning approach highlights the possible role that metacognition would play in stopping decisions. It moves beyond the threshold and reward-centric criteria that optimal stopping and foraging theories rely on by shifting the focus to psychological factors. But how can metacognition be implicated in contexts where individuals persistently persevere or prematurely stop an activity?

Quitting as a special form of stopping

Unlike simple decisions to stop an activity, for example to stop reading a book in order to eat dinner, quitting an activity, for example to stop playing a video game when things are not going well, appears to be a phenomenologically distinct category of decision, implicating a tussle between the choice to quit or not to quit before the decision is ultimately made. Typically, there is a normative expectation to 'not quit' that makes quitting decisions cognitively and emotionally challenging for people (Duke, 2022).

In this paper, we attempt to characterize the metacognitive characteristics of quitting decisions. Existing metacognitive research on stopping decisions has focused on laboratory-based stopping problems, which tend to occur on short timescales and are concerned with immediate rewards that are incurred on stopping at a particular point in the sequence (e.g. choosing a maximum out of a number sequence; (Guan, Lee, & Silva, 2014)). However, quitting is a characteristic of decisions made on larger timescales than are easily accessible in lab experiments, and thus a study of the properties of quitting is poorly matched with experimental methods. Instead, we rely on correlative analysis of a naturalistic domain with an intrinsic affordance for decision-makers to quit in repeated episodes of the same activity, rendering their behavior amenable to empirical analysis (Kuperwajs & Ma, 2022). In particular, we study the metacognitive aspects that go into chess players' decisions to resign while playing chess games online.

Chess has been popularly studied within the domains of decision making and complex planning (Charness, 1977; De Groot, 1978) by providing natural control over when a person decides to start or stop a particular game (Chowdhary, Iacopini, & Battiston, 2023; Russek, Acosta-Kane, van Opheusden, Mattar, & Griffiths, 2022). Using chess as our primary *modus operandi* for studying quitting, we focus on large-scale datasets as a testing ground to capture the importance of decision-making on human performance. Player quality can be easily characterized through player ratings defined by a standardized rating system (Glicko-2 rating, (Glickman, 2012)). Furthermore, the result of each decision can be estimated using approximate chess algorithms and compared to the actual outcome of the game (Sigman, Etchemendy, Slezak, & Cecchi, 2010).

We define 'quitting' in chess as the event of a player resigning a game without being already forced in a state of checkmate. Using this definition, we focus our efforts on mapping the game-related factors that affect when a player resigns. We assume that this action is preceded by a decision to quit the match involving latent metacognitive capabilities within the individual. Based on this formalization, we examine how chess players make decisions about when to quit a match. It is reasonable to expect that decisions of quitting would represent significant mental conflict for the player, as resigning from a match incurs a cost on a player's skill ratings.

Methods

Measure of quitting

In order to model how sequential decisions progress towards a decision to quit, survival analysis provides a convenient framework. We operationalise two interrelated concepts (Kleinbaum & Klein, 1996) for our analysis:

1. **Survival rate**: The time marker until which a subject has survived. In our dataset, this is defined by the number of moves made within the game before the game ends.

2. **Hazard rate (or quit rate)**: The probability of the event of interest happening at the next time marker, given that the subject has survived until this time marker. In our dataset, this is calculated using a combination of resign tags and last move number for all games for each player.

The quit rate (or hazard) represents a psychological variable that acts as an indicator of a player's propensity to quit a chess match. In survival analysis terms, quit rate is calculated as probability of the event-of-interest (i.e. quit) taking place at the next move given that a player have survived until the current move.

Data Collection

We extracted approximately one million online chess games belonging to the Classical game category for the year 2020 from the chess server *lichess.org*. Within this category, each player is allotted min. 30 minutes as their move times, giving them ample time to plan and make their moves. The players were randomly sampled (N_{total} = 13,000) from the data set and the matches were fed into a chess evaluation engine **(Stockfish v10)**. The engine performs a move-by-move analysis for each game and assigns a numerical value for each move, signifying its advantage w.r.t. to the player. The outcome of the game as a result of checkmate is also generated by the engine. Only games ending in a win/loss through a checkmate or resign by a player are retained. Resign games are then classified for each player's game roster. Games where the last move is not a checkmate and the game result indicates an opponent win are classified as resign games for the player. Using this criterion, we retained players (N_{resign} = 250) who had at least one game resign on their roster. All retained players had min. 5 games in their roster, in order to bring statistical reliability to our models.

Dataset

Figure 1a illustrates a typical resignation scenario, where the player (playing as white) is in a hopeless losing situation with

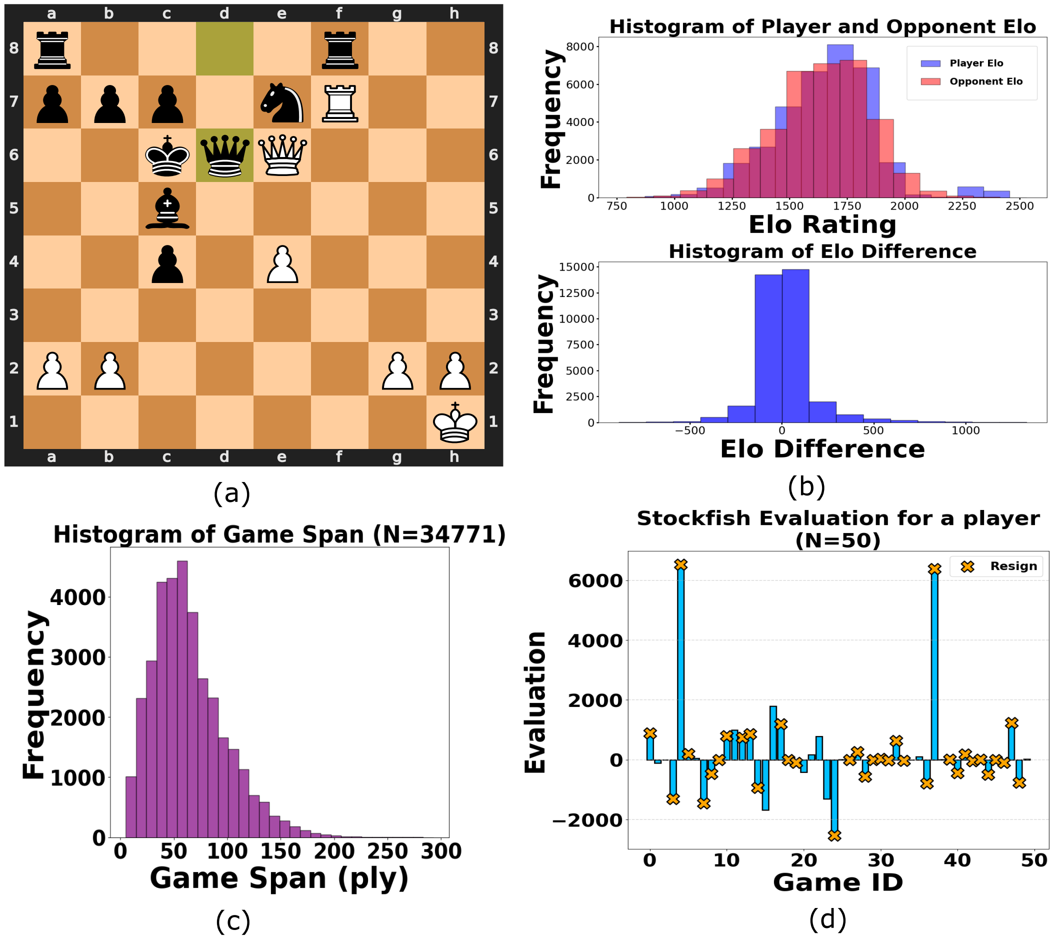

Figure 1: Summary of dataset characteristics. Panel (a) shows an example resignation scenario with respect to a white player. Panel (b) shows histograms for ELO ratings for both players and their corresponding differences. The close match indicates the ELO-matching algorithm at play in online chess queues. Panel (c) shows the distribution of game lengths for all games, in ply numbers. Panel (d) shows Stockfish evaluations for 50 games of a player. Resigned games are indicated by yellow crosses.

no chance of recovery. In our data set, the skill levels of both players were on average closely matched, as reflected in their ELO ratings, and the resulting differences clustered around zero (Figure 1b). The large sample of games offered a diverse range of game lengths, with a wide variety of scenarios (Figure 1c). We consider the complete roster of games for each player to account for the variability of decision-making between players, to capture their quitting behavior (Figure 1d).

Results

In this setup, we empirically investigated how decisions to quit are affected by game factors and past experiences. We asked two questions in our analysis: (1) How do game factors, such as skill gaps between 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?

We formulated a mixed-effects hazard regression model with game factors as fixed effects and player ID as random effects. The model captures how quit rate (i.e. the hazard rate) varies with different game factors. We included player-level game factors for the evaluation of the last move, player ELO and difference in ELO between the player and opponent, distance from the last quit match to the current match, and the number of matches that have been stopped. All these factors were indexed at the player and game level (Table 1). For the random intercept, estimated pseudo-variance was 0.086, indicating low heterogeneity in baseline hazard between players.

Game factors affect quitting decisions

Previous research has quantified human chess performance in terms of two factors in the game: the variety of opening sequences and the win rate within different skill groups (Steyvers & Benjamin, 2019; Chowdhary et al., 2023; De Marzo & Servedio, 2023). It is intuitive to expect the win rate to scale up with increasing skill. Would a similar relation hold with a player's quit rate?

We asked this question with reference to two key game fac-

Table 1: Mixed-effects Cox regression model with fixed effect for game factors and past quitting and random effect and intercept for player ID. Hazard Rate is measured as the hazard of quitting a match. Game and past quitting factors are z-score normalized.

Covariate	$exp(\beta)$	95% CI	p-value
Evaluation(last move)	0.94	0.93,0.95	<0.001
Player ELO	0.98	0.94,1.02	0.3
ELO difference	0.72	0.70,0.73	<0.001
Last Quit	0.87	0.82,0.92	<0.001
Total Quits	1.09	1.02,1.16	0.007
Evaluation × ELO Difference	0.97	0.96,0.99	0.002
Player ELO × ELO Difference	0.94	0.92,0.96	<0.001
Last Quit × Total Quits	0.93	0.88,0.99	0.021
ELO Difference × Last Quit	0.99	0.97,1.01	0.3
Evaluation(last move) × Last Quit × Total Quits	1.03	1.01,1.04	<0.001

tors: evaluation of the last move in the game and the skill difference between a player and their opponent. The last move's evaluation provided by Stockfish is assumed to represent the 'ground truth' of the board state that a player may approximate and evaluate in their decision-making process (Holdaway & Vul, 2021). For quantifying skill level, we take skill difference as the metric instead of the absolute ELO values for the player, as we believe that the difference represents a normalized metric for a player to judge their confidence of winning/losing in a match.

Table 1 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 the difference in ELO ($\exp(\beta) = 0.72$, 95%CI: 0.70-0.73, $p < 0.001$) and the evaluation of the last move ($\exp(\beta) = 0.94$, 95%CI: 0.93-0.95, $p < 0.001$) predicted a 28% and 6% decrease in the risk of quitting a match. This indicates that a player's decision to quit is influenced by different game factors somewhat counterintuitively: playing against a significantly superior player strongly predisposes players to persist and having made a significantly poor move produces a directionally similar but weaker effect. Interestingly, the interaction between last move and ELO difference ($\exp(\beta) = 0.97$, 95%CI: 0.96-0.99, $p < 0.002$) predicts a 3% decrease in the risk of quitting. The effect of the difference in ELO persists even after adjusting for player skill ($\exp(\beta) = 0.94$, 95% CI: 0.92-0.96, $p < 0.001$) through a 6% decrease in the hazard of quitting. The latter interaction could also be due to covariance between a player's ratings and the difference with regard to the opponent, which is factored into the Glicko-2 rating scheme.

Quitting often makes quitting more likely

In order to check if past experiences of quitting affect the player's current quitting decision, we defined two novel variables: "Last Quit" defines how far in time the current match is from the most recently quit match. "Total Quits" refers to the total matches the player has quit until the current match. The Last Quit factor ($\exp(\beta) = 0.87$, 95%CI: 0.82-0.92, $p < 0.001$) predicted a 13% decrease in the hazard of quitting, suggesting that the experience of having recently quit makes quitting in the current game less likely, keeping other factors constant. In contrast, Total Quits ($\exp(\beta) = 1.09$, 95%CI: 1.02-1.16, $p = 0.007$)) predicted a 9% increase in the hazard of quitting. This effect reverses after adjusting for the interaction with Last Quit ($\exp(\beta) = 0.93$, 95%CI: 0.88-0.99, $p = 0.021$)), predicting a 7% decrease, although the statistical significance of this interaction does not survive corrections for multiple comparisons at $p = 0.05$.

The quality of the board state interacts with the past quits in time and frequency ($\exp(\beta) = 1.03$, 95%CI: 1.01-1.04, $p < 0.001$)), increasing the hazard of quitting by 3%. This indicates a complex interaction between current board quality and previous quitting experiences on the decision to quit, likely driven by greater player experience.

Tilting towards a decision to quit

The complex nature of the relationship between game factors and past quitting experiences toward predicting the probability of quitting should also translate into subsequent behavior. A pervasive aspect of playing games has been the influence of negative game performance on decision making (Schüll, 2016) and the emotions of the player. This tilting behavior has been documented in poker players engaged in loss chasing (B. R. Browne, 1989) and also in esports players (Wu, Lee, & Steinkuehler, 2021). A behavioral consequence of quitting decisions could be how fast a player starts the next match after having quit relative to a match where they have not made this decision. Players may 'tilt' towards quickly starting the next match after a quit in order to recover the loss and embarrassment of resigning. So, we were interested in checking if the player's time to the next game is affected by these quitting decisions and game factors.

We measured the hunt for new games as the absolute deviation in time between a new game and resign game. We normalized it with respect to the average deviation in time to play the next game for each player. This indicates how much the time to the next game (after a resign) deviates from the mean time for each player.

Table 2: Mixed-effects linear model for the time to play the next game. Differences in ELO ratings and time since last quit are included as fixed-effects with player ID for random intercepts.

	Coef. β	Std. Error	95% CI	p-value
Intercept	18.976	1.353	16.224,21.529	<0.0001
ELO difference	0.257	0.138	-0.013,0.528	0.062
Last Quit	0.267	0.124	0.025,0.509	0.030

We formalized a mixed-effects linear regression model of the form: Time to next game $\sim$ ELO difference + Last Quit + $1|$player. The model follows a random-effects structure with time to the next game as the dependent variable. The difference in ELO and the time since the last quit are included as predictors in the model. These predictors were z-score normalized for the analysis. The player ID represents the random-effects variable in the model.

Table 2 shows the results of the mixed effects model. The main effect of time since the last quit ($\beta = 0.267, SE = 0.124, p = 0.030$) predicted the time to play the next game. If a player had a recent quit, they are more likely to play the next game instead of waiting. This demonstrates a 'tilting' behavior for the player. However, no main effect of the ELO difference was observed ($p = 0.062$). Skill differences have no influence on when the player starts their next match, but their past decisions about quitting do. This indicates a mismatch between the decision making during a match and its downstream consequences on future play when a player decides to quit.

For this mixed linear model, following Nakagawa and Schielzeth (2013), the marginal R^2 was **0.05**, indicating that fixed effects explained 5% of the variance in deviation in time to next game after a resign. The conditional R^2 was **0.56**, indicating that 56% of the variance was explained when accounting for both fixed and random effects. Both homoscedasticity and normality assumptions were violated in our data. Possible reasons could be the unequal sizes of data available for each player, which would affect the variance of error residuals for the random-effects component of the model. Despite these violations, model estimates can be interpreted as robust with minimal consequences (Schielzeth et al., 2020).

Discussion

Deciding to quit a particular activity is dependent on both task-specific and task-invariant psychological factors. Previous research has focused extensively on understanding when people stop seeking information to make a decision to quit or persist (Guan et al., 2014; Kuperwajs & Ma, 2022; Sukhov, Dubey, Duke, & Griffiths, 2023). However, due to the abstract nature of stopping decisions in laboratory tasks, the metacognitive covariates involved are poorly understood. In this paper, we sought to identify metacognitive covariates for a real-world quitting decision, when to quit a game of chess, with the help of a large-scale database of online chess.

We found that larger skill differences between a player and their opponent seem to push them to persevere in the match. Holdaway and Vul (2021) obtained a similar relation with risk preferences in players. They showed that players may prefer to play riskier moves against opponents with greater skill. They postulated that this could be a strategy used by players to outsmart the opponent, especially in losing situations. This is further solidified by a game board state that is disadvantageous to the player, as shown in our analysis. Thus, players are, somewhat counterintuitively, less likely to quit in adverse game situations, and when paired against a stronger opponent. The simplest explanation for this pattern of behavior is that players expect to be doing poorly when playing against stronger opponents and thus are less conflicted metacognitively about the choice to continue playing.

We also find evidence of behavior consistent with a sense of guilt over violation of normative expectations, with players less likely to quit a game if they have recently quit a game, but more likely to quit a game if they have a history of quitting games (and thus habituated to norm violation). Loss chase, that is, being impelled to play again sooner after having quit in a game, is also consistent with this sense.

Metacognitive failure of will

A complete account of metacognition requires that both the monitoring and control processes work together. Our results indicate a deficit in the control process. The players monitor the current state of the game, the difference in skill, and the history of quitting. But when they have to implement this rational choice into action, they seem to fail. This is illustrated by the time it takes for them to start their next match. If they were aware of their history of quitting, they would act in a way that prevents more losses to their ELO ratings. Pursuing a win during a cold streak would be more difficult to achieve compared to a hot streak. Although players do encounter longer durations of cold streaks compared to hot streaks (Chowdhary et al., 2023), it seems almost irrational that a player should continue onward to their next game right after a match quit. This discrepancy is highlighted by our analysis and could possibly hint at the lack of metacognitive control in such instances. According to at least one classic theory of metacognition (Nelson, 1990), metacognitive control involves planning w.r.t. retrieval of relevant past experience on judgments about task performance. Based on our results of tilting in chess games, we postulate the cause of this experience as a metacognitive failure that affects willpower.

How would we begin to model this breakdown of the will? Figure 2 illustrates a network model of willpower. Each node j represents a quanta of willpower capacity measured as $W_j = (1+\alpha)L_j$ (Figure 2a). This willpower capacity is proportional to the load carried by the node, defined by its degree of connection to neighbor nodes. Thus, each node is

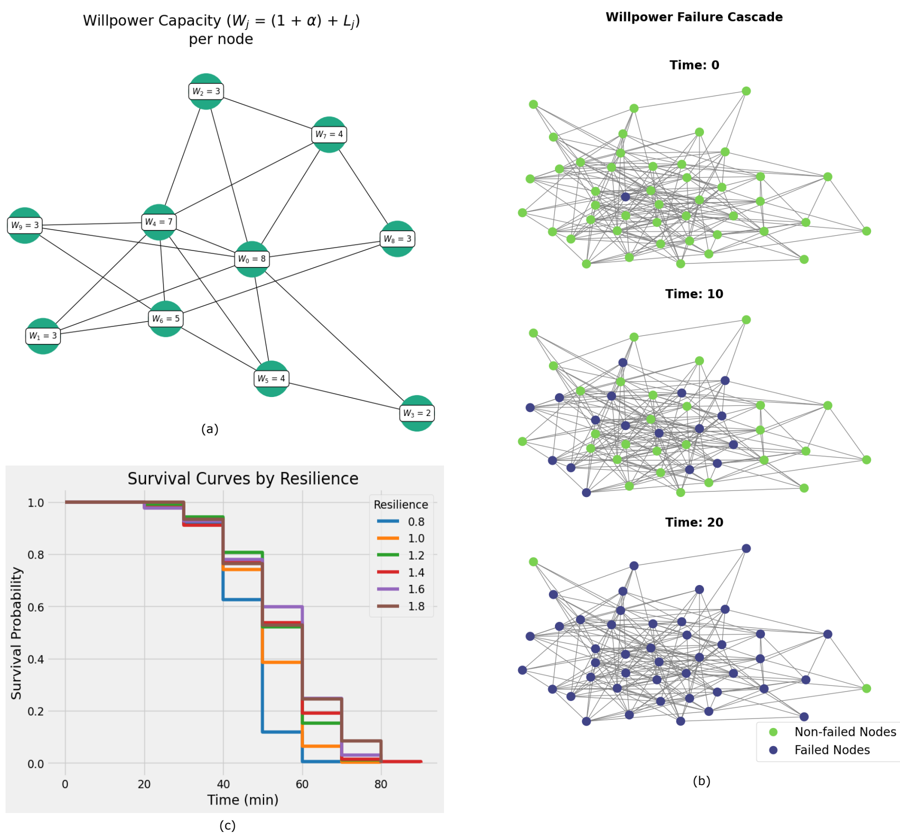

Figure 2: A conceptual network model of willpower failure. Panel (a) shows an example network with willpower per node. Willpower (W_j) is proportional to load (L_j) for each node. Panel (b) shows a graph with cascade failure of willpower across time. Failure iniated by random node. Panel (c) shows a survival curve against willpower failure cascade for individuals with varying resilience. For Panels (a) and (c), a Barabasi-Albert graph is used with 10 nodes (3 edges per node) and 5000 nodes (4 edges per node), respectively. For Panel (c), an Erdos-Renyi graph with 50 nodes and probability $p = 0.2$ for growing each edge.

supported by its neighbor in keeping the network connected.

The distinct phenomenology of quitting decisions is influenced by the intrusive thoughts of quitting that spontaneously appear in a person's awareness. A single onset of a quitting thought can initiate an autocatalytic effect on the breakdown of a person's will to persist on the task. Figure 2b illustrates the cascading effect of a single quitting thought. Failure of a random node can initiate downstream consequences of a network collapse, over time. This type of cascade failure would result in a rapid breakdown of network integrity (Motter & Lai, 2002).

What protects against the tyranny of these quitting thoughts? Resilience, the ability to respond to a task after failure (Werner, 1989), has been widely cited as a trait that individuals possess in adverse situations. These individuals tend to show persistent behaviors on tasks where the cost of persisting tends to be greater than its benefit. Would resilience protect the individual against weathering the quitting storm? Figure 2c describes an example case of survival against cascade failure of willpower for individuals with different levels of resilience, operationalized as the variable α. This "battle of wills" can be modulated by optimal metacognitive control, which could explain the variations in levels of resilience demonstrated by an individual in different contexts (Matthews, Panganiban, Wells, Wohleber, & Reinerman-Jones, 2019).

Thus, rather than attribute willpower as a trait (Ainslie, 2021), as is conventional, we propose that it may be insightful to study it as an outcome of metacognitive control. Future work can characterize how components of this network model could operationalize psychological factors that affect quitting decisions in task contexts. The prediction of the network model of an autocatalytic pattern for quitting thoughts can also be studied in additional ecologically rich domains (Hutchins, 1995). We expect that further pushing the envelope of ecological embeddedness to study metacognitive behavior, as demonstrated in this work, will yield decision-making models with greater phenomenological validity.

References

Ainslie, G. (2021). Willpower with and without effort. *Behavioral and Brain Sciences*, *44*, e30.

Branco, F., Sun, M., & Villas-Boas, J. M. (2012). Optimal search for product information. *Management Science*, *58*(11), 2037–2056. Retrieved from https://doi.org/10.1287/mnsc.1120.1534

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. Retrieved from https://www.jstor.org/stable/25148782

Bugbee, E., & Gonzalez, C. (2022). Deciding when to stop: Cognitive models of sequential decisions in optimal stopping tasks. *In preparation*.

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. *Chess skill in man and machine*, 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.

De Groot, A. D. (1978). *Thought and choice in chess* (Vol. 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.

Glickman, M. E. (2012). Example of the glicko-2 system. *Boston University*, *28*.

Green, R. F. (1984). Stopping rules for optimal foragers. *The American Naturalist*, *123*(1), 30–43. Retrieved from https://www.jstor.org/stable/2461356

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* (Vol. 36).

Guan, M., Lee, M. D., & Vandekerckhove, J. (2015). A hierarchical cognitive threshold model of human decision making on different length optimal stopping problems. *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, 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* (Vol. 43).

Hutchins, E. (1995). *Cognition in the wild*. MIT press.

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

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* (Vol. 44).

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

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

Matthews, G., Panganiban, A. R., Wells, A., Wohleber, R. W., & Reinerman-Jones, L. E. (2019). Metacognition, hardiness, and grit as resilience factors in unmanned aerial systems (uas) operations: a simulation study. *Frontiers in psychology*, *10*, 640.

Motter, A. E., & Lai, Y.-C. (2002). Cascade-based attacks on complex networks. *Physical Review E*, *66*(6), 065102.

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 *Psychology of learning and motivation* (Vol. 26, pp. 125–173). Elsevier.

Russek, E., Acosta-Kane, D., van Opheusden, B., Mattar, M. G., & Griffiths, T. (2022). Time spent thinking in online chess reflects the value of computation.

Schielzeth, H., Dingemanse, N. J., Nakagawa, S., Westneat, D. F., Allegue, H., Teplitsky, C., … 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. doi: 10.1111/2041-210X.13434

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

Sigman, M., Etchemendy, P., Slezak, D. F., & Cecchi, G. A. (2010). Response time distributions in rapid chess: a large-scale decision making experiment. *Frontiers in Decision Neuroscience*, *4*, 60.

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*.

Werner, E. E. (1989). Vulnerability and resiliency: A longitudinal perspective. *Children at risk: Assessment, longitudinal research, and intervention*, 158–172.

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. Retrieved 2025-01-30, from https://dl.acm.org/doi/10.1145/3411764.3445143 doi: 10.1145/3411764.3445143