Exploring resource-rational planning under time pressure in online chess

Ionatan Kuperwajs (ikuperwajs@princeton.edu)
Department of Computer Science, Princeton University
Princeton, NJ, United States

Evan M. Russek (erussek@princeton.edu)
Department of Computer Science, Princeton University
Princeton, NJ, United States

Lisa Schut (schut@robots.ox.ac.uk)
Department of Computer Science, University of Oxford
Oxford, United Kingdom

Yotam Sagiv (yotam.sagiv@nyu.edu)
Department of Psychology, New York University
New York, NY, United States

Marcelo G. Mattar (marcelo.mattar@nyu.edu)
Center for Neural Science, New York University
Department of Psychology, New York University
New York, NY, United States

Wei Ji Ma (weijima@nyu.edu)
Center for Neural Science, New York University
Department of Psychology, New York University
New York, NY, United States

Thomas L. Griffiths (tomg@princeton.edu)
Department of Computer Science, Princeton University
Department of Psychology, Princeton University
Princeton, NJ, United States

Abstract

Human planning is incredibly efficient. Even in complex situations with many possible courses of action, people are able to make good decisions. Recent proposals suggest that a primary contributor to this efficiency is the intelligent use of cognitive resources, but how people allocate these resources under time constraints is not fully understood. In this work, we conduct a resource-rational analysis of planning in a large data set of online chess games. We first demonstrate that players spent more time thinking when they had more time to do so, and that this effect was especially prevalent when computation was more valuable. Then, we show that additional time spent planning resulted in better selected moves when one existed, and compare between signals of general and immediate time pressure. Finally, we highlight the role of expertise in this setting. Our results provide evidence that people make resource-rational choices when planning under time pressure.

Keywords: decision-making; planning; resource rationality; time pressure

Introduction

Decision-making in the real world often involves considering many different courses of action and their consequences, a process known as planning. A key characteristic of human cognition is the ability to plan efficiently (Miller & Venditto, 2021; Mattar & Lengyel, 2022). Faced with a computationally intractable number of alternatives to evaluate, people are still able to make good decisions within a constrained time budget. This is particularly evident in combinatorial games such as chess (de Groot, 2014).

Recent work has taken the view that cognitive mechanisms underlying processes like planning can be revealed by considering how to make optimal use of limited resources (Griffiths, Lieder, & Goodman, 2015; Gershman, Horvitz, & Tenenbaum, 2015). This approach of resource-rational analysis strives towards optimality by deriving models of human behavior that take into account which cognitive operations are available to people, how long they take, and how costly they are. Assuming that people select actions using mental strategies that strike a balance between the utility of the chosen action and the cognitive cost of making the decision generates predictions about how people will navigate such tradeoffs. Resource-rational analysis has advanced our understanding of how individuals allocate cognitive resources, and has shown that they are sensitive to the costs and benefits of executing computations (Kool, Gershman, & Cushman, 2017; Lieder & Griffiths, 2020; Frömer, Lin, Dean Wolf, Inzlicht, & Shenhav, 2021). Applied to planning, this hypothesis has been tested by measuring how much time humans spent thinking before acting in a large-scale online chess data set (Russek, Acosta-Kane, van Opheusden, Mattar, & Griffiths, 2022). Players spent more time thinking in board positions where additional computation was more beneficial, and this relationship was greater in stronger players.

A related but still open question is how people plan given different time constraints. An attempt to answer this question was made in a combinatorial game called 4-in-a-row that is simpler than chess but more complex than the majority of tasks used in cognitive science (van Opheusden & Ma, 2019). In this game, strong play requires thinking multiple steps ahead, but the task is simple enough that it remains amenable to process-level modeling. By fitting a computational model to human decisions, van Opheusden et al. (2023) uncovered robust evidence for increased planning depth with expertise. In a time pressure experiment where players were given differing time limits per move, planning depth scaled with each condition. However, no improvement in the participants' playing strength was found, potentially due to correspondingly increasing attentional oversights. Historically, time pressure has also been studied in the context of chess, where manipulations that aim to selectively impair search while leaving pattern recognition abilities intact seem to affect experts more than novices (Holding, 1992).

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

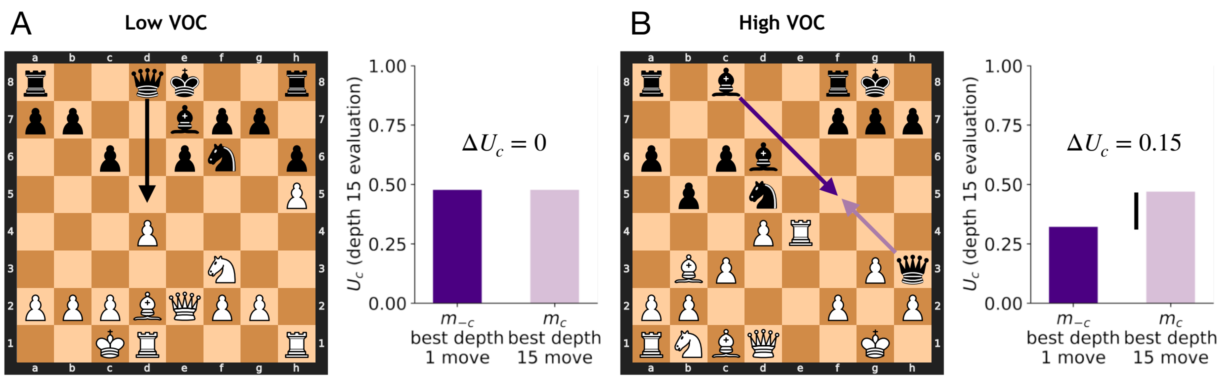

Figure 1: Value of computation. **(A)** Board position with a low VOC (left). The move `Qd8-d5` maximizes both Stockfish's depth 1 and depth 15 evaluations, and thus the optimal move with no planning (m_{-c}) and the optimal move with extensive planning (m_c) are the same. Since computation does not cause the selected move to change, it provides no benefit (right). **(B)** Board position with a high VOC (left). At depth 1, the move `Bc8-f5` maximizes Stockfish's evaluation. However, a depth 15 search reveals that the move `Qh3-f5` results in a more favorable board position advantage (U_c), measured as the estimated probability of winning from this position ignoring time remaining. Since this computation finds a better selected move, it provides a benefit of approximately 0.15 (right).

Here, we merge these lines of research to provide a resource-rational account of human planning under time constraints. We use chess as a setting to conduct this analysis, in part to extend previous work but also due to the complexity that chess positions present to the decision-maker. From a researcher perspective, chess provides the opportunity to explore the interaction between time pressure and expertise. The popularity of online chess platforms, which have resulted in massive behavioral data sets, and modern chess engines, which provide better tools for characterizing human decision-making and planning, enable our analysis. We leverage these two developments to estimate the benefits of applying planning computations across different positions, selecting two frequently occurring positions to investigate further. Disentangling the various factors influencing a complex process such as planning requires identifying repeating situations that can be investigated in depth. Our data set facilitates this empirical approach embedded in a naturalistic environment. We demonstrate that move times and move quality are influenced by time pressure, and that these effects are modulated by the value of computation in each position. Finally, we highlight the relationship between expertise, measured by playing strength, with time spent planning across different time control settings.

Methods

Our approach to characterizing the effects of time pressure on planning in chess is based on prior work (Russek et al., 2022; Kuperwajs, van Opheusden, Russek, & Griffiths, 2024). Specifically, we make use of the open-source chess engine Stockfish (`https://stockfishchess.org`) as well as a large-scale data set of human decisions from the online chess platform Lichess (`https://lichess.org`). In the following sections, we outline the role that each of these components play in our investigation.

Move evaluation

A central aim of our analysis is to uncover the effects of time pressure in situations with varying costs and benefits. Stockfish provides us with an idealized model of computation in order to do so. Like other modern chess engines (Campbell, Hoane Jr, & Hsu, 2002; Silver et al., 2018) and models of human planning in combinatorial games (van Opheusden et al., 2023), Stockfish employs heuristic search to estimate a player's advantage from a given board position. This is done in two parts: a static function that uses a neural network trained through supervised learning to quickly map board states to approximate evaluations, and an iterative deepening tree search that improves those evaluations by searching board states over likely future positions.

With this tool, we can approximate the value of any move in any board position in terms of the estimated probability of winning from the resultant position. By repeating this over many candidate moves, we can compute the optimal move at any specified depth of search. We can thus compare how move preferences shift with additional computational resources. Following Russek et al. (2022), we abstract away the particulars of node-to-node choices about where to search and coarsely compare the effects of conducting a depth 1 and a depth 15 search. Practically speaking, Stockfish serves as a ground truth for evaluating actions in our data set. We used Stockfish version 14 for all evaluations, and interfaced with Stockfish through the Python chess package.

Value of computation

Rational metareasoning indicates that the decision of when to spend resources planning should be based on a cost-benefit comparison (Russell & Wefald, 1991; Horvitz, 1991). In

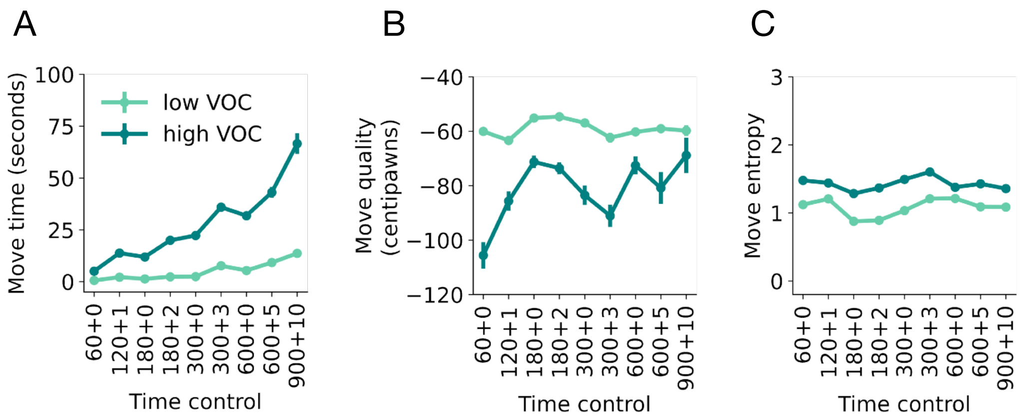

Figure 2: Resource-rational behavior under time pressure. **(A)** Time taken to make a move measured in seconds as a function of time control setting for the low and high VOC positions. Here and in all subsequent panels, the data are displayed as the mean across all players, with error bars indicating the standard error of the mean if appropriate. **(B)** Quality of the move made by the active player measured in centipawns as a function of time control setting for the low and high VOC positions. Negative units of centipawns indicate a losing position for the active player. **(C)** Entropy of the distribution over all moves made by active players as a function of time control setting for the low and high VOC positions.

chess, the benefit of planning is that it can lead a player to make a move that results in a more favorable position and thus increase the chances of winning. The cost of planning is an opportunity cost, as spending time planning leaves less time available to consider future moves. Using this framing, we define the value of computation (VOC) in a given position. Here we provide sufficient details for understanding VOC in the context of this work, but see Russek et al. (2022) for further details. We assume that players start each turn with an initial estimate of the utility of candidate moves, U_{-c}, which involves no planning. Utility provides a measure of the player's board position advantage, where $U = 1$ is a certain win and $U = 0$ is a certain loss if the optimal move is played. Players can then make the maximum utility move that is available to them m_{-c}, or perform a planning computation. We assume that planning can provide the true utility of moves U_c, allowing players to select the optimal move m_c. Thus, VOC is defined as the difference in board position advantage between these two moves:

$$\Delta U_c = U_c(m_c) - U_c(m_{-c}). \tag{1}$$

Intuitively, VOC is high when planning changes the preferred move and improves the resultant board position relative to the previously preferred move. VOC is 0 when planning does not change the preferred move. In our analysis, we operationalize m_{-c} as the move that Stockfish would select following a depth 1 search and m_c as the move that Stockfish would select following a depth 15 search. Note that at this depth, Stockfish plays above grandmaster level (Ferreira, 2013).

Data

We use a data set consisting of games played on the online chess platform Lichess. Each game contains the sequence of moves that were made by each player. Games are played across a variety of time control settings which denote the amount of time that players start with on their clock and an increment that is added back to the that clock after each move. Players lose when they run out of clock time, enabling our investigation of time pressure. Each player also maintains an Elo rating that reflects their overall playing strength relative to other players (Elo, 1978). Higher Elo indicates a stronger player, and this rating is updated after each game depending on the outcome.

To study the effects of time pressure, we wanted to find board positions that repeat often in the data set across each time control setting and avoid positions that occur in the opening or ending of games as these can involve memorized sequences of moves or abnormal play rather than planning. To select positions for further analysis, we counted the frequency of board positions from all games played in 2022 and 2023 that occurred after the 30th move ply and had at least 10 pieces on board. We then computed an associated VOC for these positions. Stockfish returns evaluations in units of either centipawns or distance that indicate the number of move plys from mate. In computing VOC, to place these separate units on a common scale, we converted each of these units to a new measure of position advantage by fitting two logistic regression models which mapped either centipawn advantage or distance from mate to the probability that the active player won the match. To select m_c, we also control for the fact that, due to the pruning heuristics that Stockfish employs, the best move is not always identified at depth 15. We thus select m_c as the move with maximum value, out of a consideration set consisting of each move Stockfish selects at each depth and computing U_c for each move in this set.

After this process, we selected one of the most frequently

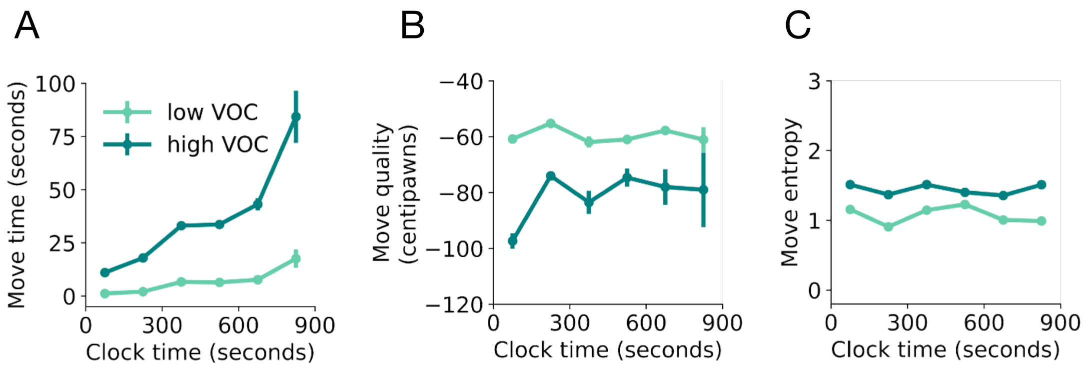

Figure 3: Resource-rational use of clock time. **(A)** Time taken to make a move measured in seconds as a function of the remaining clock time for the low and high VOC positions. Here and in all subsequent panels, results are binned into 6 bins and the data are displayed as the mean across all players, with error bars indicating the standard error of the mean if appropriate. **(B)** Quality of the move made by the active player measured in centipawns as a function of the remaining clock time for the low and high VOC positions. Negative units of centipawns indicate a losing position for the active player. **(C)** Entropy of the distribution over all moves made by active players as a function of the remaining clock time for the low and high VOC positions.

occurring positions with a VOC of 0 (Figure 1A) and another of the most frequently occurring positions with a relatively high VOC of approximately 0.15 (Figure 1B). This resulted in $12,384$ moves in the former case and $12,786$ in the latter case, distributed across 9 time control settings that each had at least 100 moves per position. We note that, in contrast to Russek et al. (2022), we intentionally control for board position to characterize how use of time affects behavior. Time control settings use the notation $S+I$ which indicates that each player starts with S seconds on their clock and gets back I seconds following each move as an increment. The conditions included in our analysis that are treated as independent categories are $60+0$, $120+1$, $180+0$, $180+2$, $300+0$, $600+0$, $600+5$, and $900+10$. These make up 9 of the 11 default time control settings on Lichess, and we excluded $1800+0$ and $1800+20$ due to a lack of data.

Results

Given a large set of moves made by chess players in a repeated low and high VOC board position, we now examine how time pressure affects decision-making in these positions, and how this changes with expertise.

Rational use of resources under time constraints

We hypothesized that time control settings might have an effect on three signatures of decision-making in chess: move time, move quality, and move entropy. Additionally, we expected these relationships to vary across the two positions we analyzed. We computed move times as the difference in clock time between a player's successive moves, minus the time control specified increment. To measure move quality, we once again used Stockfish to evaluate the move a player made at a depth 15 search. Finally, we defined entropy H over the distribution of moves made by all players p_n such that:

$$H = -\sum p_n \log(p_n). \tag{2}$$

Using these metrics, we found evidence of resource-rational behavior. As players had more time to play a game, they spent longer thinking before making a move (Figure 2A). This was modulated by the potential benefits of computation, as move times increased much more rapidly as the time control setting became more relaxed in the high VOC as opposed to the low VOC position. This indicates that players spent more time thinking, and therefore planning, when presented with the opportunity to spend more time in situations that would reward doing so. This raises the question: did players actually find better moves as the time they spent thinking increased? This was indeed the case, but only in the high VOC position where there was the potential for move quality to increase in the first place (Figure 2B). The increase in move quality as a function of time control setting plateaued after an initial increase, suggesting that perhaps there is a minimum amount of time that is necessary in order to select the best move. In the low VOC position, there was no effect of time control setting on move quality. Finally, move entropy was constant across all time control settings irrespective of VOC (Figure 2C). As such, more allotted time did not reduce the variance of moves selected by the population of players in our data set.

Time control setting is ultimately a general indicator of time pressure, as players select the setting they play under and have that in mind as they make moves throughout the game. In order to test if these results changed under immediate time pressure, we repeated the same analysis using clock time. Since clock time changes as players spend time thinking about a move and can result in losing the game altogether, it has the potential to provide a more direct signal of time pres-

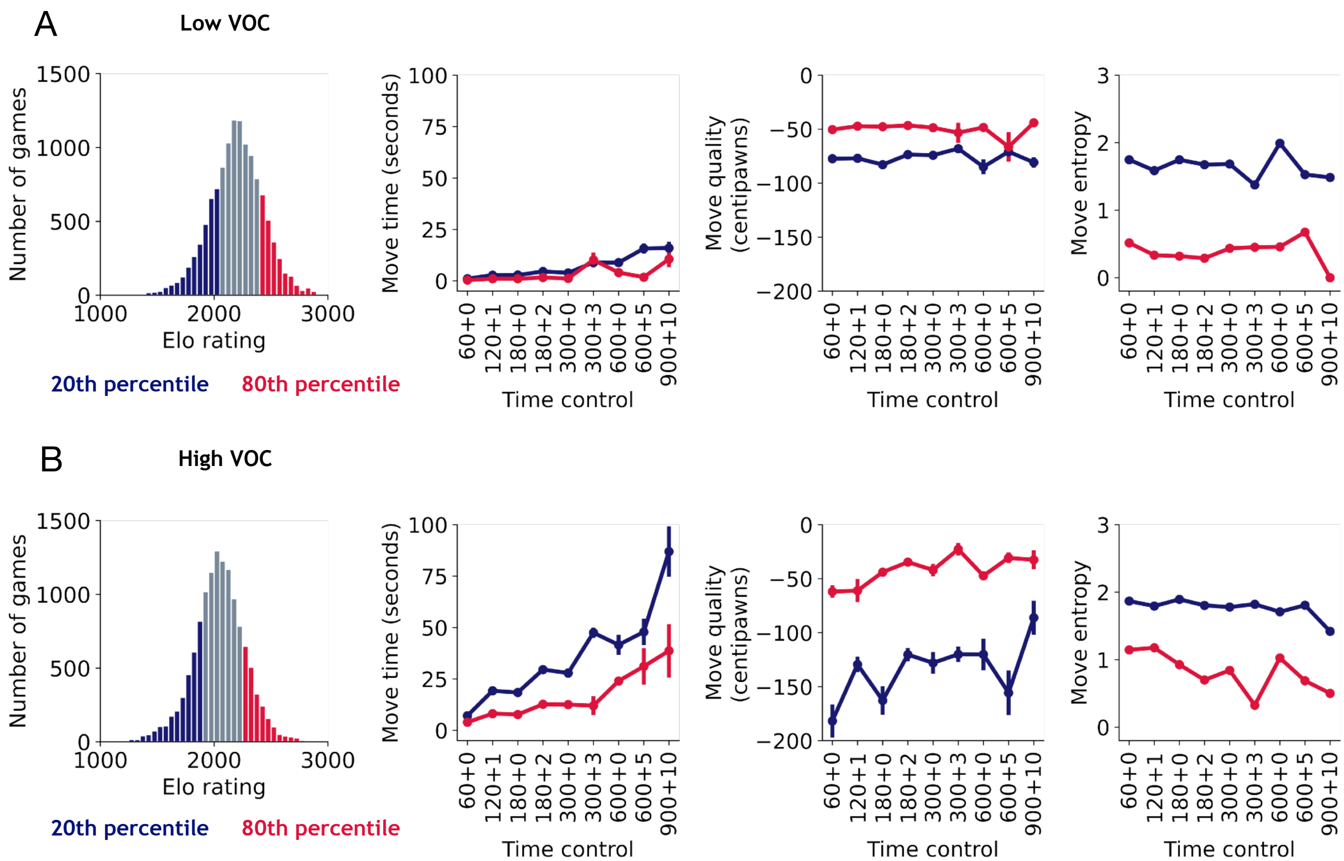

Figure 4: Expertise modulates the effect of time pressure. **(A)** Distribution of playing strength measured as an Elo rating over all players in the low VOC position. The following plots replicate those in Figure 2, but split by the players in the 20th and 80th percentile of Elo ratings rather than by VOC. **(B)** Same as in (A), but for all players in the high VOC position.

sure. We found roughly the same effects of clock time on move time, move quality, and move entropy, including how these effects varied with VOC (Figure 3). The one exception was time spent thinking in the high VOC position, which was longer overall. This suggests that clock time and time control setting largely have the same influence in driving resource-rational behavior, but that more time on the clock signals to players that they can spend more time selecting a move compared to simply knowing what condition they are in.

The effects of expertise on planning

If modulation of move time relative to time pressure is a critical factor for efficient planning, we would expect this relationship to be sensitive to expertise. To examine this in the data, we tested whether our metrics varied with player Elo rating. We compared between the subset of weakest and strongest players in each VOC position, taking the 20th and 80th percentile of players respectively. In the low VOC setting, there was a minimal impact of time control setting on move time, move quality, and move entropy (Figure 4A). The main visible effect is of expertise: experts spent less time thinking, made better moves, and had less variance in the distribution of moves they made compared to novices. The intuition here is that players with a higher Elo rating are overall better in terms of their chess abilities, and as such need to spend less time thinking in order to arrive at consistently good moves regardless of the amount of time available to them.

By contrast, the high VOC position revealed a significant interaction between time control setting, playing strength, and move time, move quality, and move entropy (Figure 4B). As expected, players spent more time thinking when they had more time available. This was especially true of novice players, who likely needed the extra time to plan compared to better players. This was validated by move quality, which drastically increased for the weaker players as a function of time control setting. Stronger players also found better moves with more time, but already made quite good moves in the most strict time control setting. Interestingly, we did find an effect of time pressure on move entropy in this analysis. This was particularly true in the case of stronger players, who converged on similar moves with more time to plan. Note that in the high VOC position, the general differences between experts and novices present in the low VOC position were preserved.

Discussion

In this work, we analyzed whether people made decisions consistent with resource rationality under time pressure. We did this by first identifying frequently occurring chess positions in which players might need to rely on mental simulation as opposed to memorization or other cognitive processes in order to select a good move. Then, we compared players' choices when the value of computation, or the benefit of planning as measured by the increase in board position advantage calculated by a chess engine, was either 0 or relatively high. We found that thinking times increased as players had more time to plan, and that this effect was more pronounced in the position with higher value of computation where players also made better moves. We replicated these results using clock time instead of time control setting in order to test a signal for immediate as opposed to general time pressure. Finally, we investigated the role of expertise in these decisions, highlighting that playing strength modulates move time, move quality, and move entropy, but only when there was a benefit to thinking ahead. Taken together, our results demonstrate that planning in complex environments such as chess shows signatures of sensitivity to the benefits and costs of computation under time constraints, and that this sensitivity is dependent on expertise.

While our analysis makes use of a large-scale chess data set to identify positions that enable planning to occur, a limitation is that we highlight two specific states across all of our results. In future work, we plan to extend this to a wide range of board positions that tile the entire VOC space such that we capture subtle features of those positions that influence decision-making. Additionally, we aim to categorize board positions by different metrics, such as complexity estimated by the number of nodes that need to be searched before the best move is encountered or volatility estimated by the number of times the principal variation changes during search, to study how our results generalize. We can also begin to ask novel questions using our approach. For instance, does the opportunity to search lead to more learning, and how does the amount of training data given to individuals, measured by the number of games they play, affect performance?

To properly contextualize our findings, we can return to the two lines of research we set out to study simultaneously. In terms of resource-rational analysis, we primarily extended previous work on the benefit of computation in chess that applies a similar approach to focus on the added effect of time pressure (Russek et al., 2022). Related studies have found relationships with time remaining on a player's clock, the difference in value between the best and second best move, and position complexity (Sunde, Zegners, & Strittmatter, 2022) as well as trained feature-based estimators to predict move times (Burduli & Wu, 2023). Our work is consistent with these results. More broadly, a collection of experiments have shown that human deployment of cognitive control and mental effort is sensitive to variations in the rewards, costs, and efficacy of control deployment (Hall-McMaster, Muhle-Karbe, Myers, & Stokes, 2019; Grahek, Frömer, Prater Fahey, & Shenhav, 2023). In these settings, the opportunity cost of spending time is framed according to the average reward rate. By examining the consequences of resource rationality in a more complex task like chess, we confirm that this framing continues to be informative when people's goals differ from maximizing reward.

We can reach a similar conclusion with regards to the planning literature, as a number of theories have been proposed which imply that the decision of when to plan should be sensitive to principles of cost-benefit computation (Sezener, Dezfouli, & Keramati, 2019; Mastrogiuseppe & Moreno-Bote, 2022). One difference emerges when comparing with the time pressure experiment in 4-in-a-row (van Opheusden et al., 2023), which found that people plan more when given more time per move but do not improve in terms of their playing strength. In our case, we found that people did make better moves in more relaxed time conditions. Reconciling this distinction is an important problem that our current analysis does not address, but a possibility is that players in our data set did not experience the same level of attentional oversights in relaxed time conditions because they were highly motivated to win. To test this, we could look at blunder rates in known positions as a measure of errors that players make when they are not paying enough attention. An alternative explanation is that move quality is more sensitive to time constraints than Elo ratings are despite the two being correlated. Regardless, we view our work as a first step towards more deeply understanding the effects of time pressure on planning.

Finally, although our findings provide evidence for human behavior in chess that is resource-rational in relation to the benefits of computation and time constraints, we do not address the underlying process by which people produce such behavior. In fact, developing process-level theories of human planning that accurately predict the moves of individual chess players has proven to be difficult (Gobet & Jansen, 1994; Gobet, 1997). Constructing such a model could use the framework of heuristic search (Pearl, 1984; Bonet & Geffner, 1999) as a foundation to then include additional components of planning models (Huys et al., 2012; Callaway et al., 2022; Kuperwajs, Ho, & Ma, 2024) as well as knowledge specific to chess (Chase & Simon, 1973; Campitelli & Gobet, 2004). With this class of model, we can use derived model parameters and predictions to gain insight into human decision-making. Large-scale data in tandem with modern computational methods will play a central role in the development of this work.

Acknowledgments

This work was made possible by grants from the National Science Foundation (number IIS-2312373) and the NOMIS Foundation.

References

Bonet, B., & Geffner, H. (1999). Planning as heuristic search: New results. *European Conference on Planning*, 360–372.

Burduli, G., & Wu, J. (2023). Time management in a chess game through machine learning. *International Journal of Parallel, Emergent and Distributed Systems*, *38*(1), 14–34.

Callaway, F., van Opheusden, B., Gul, S., Das, P., Krueger, P. M., Griffiths, T. L., & Lieder, F. (2022). Rational use of cognitive resources in human planning. *Nature Human Behaviour*, *6*(8), 1112–1125.

Campbell, M., Hoane Jr, A. J., & Hsu, F.-h. (2002). Deep blue. *Artificial Intelligence*, *134*(1-2), 57–83.

Campitelli, G., & Gobet, F. (2004). Adaptive expert decision making: Skilled chess players search more and deeper. *ICGA Journal*, *27*(4), 209–216.

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, *4*(1), 55–81.

de Groot, A. D. (2014). *Thought and choice in chess* (Vol. 4). Walter de Gruyter GmbH & Co KG.

Elo, A. E. (1978). *The rating of chess players, past and present*. Arco Pub.

Ferreira, D. R. (2013). The impact of the search depth on chess playing strength. *ICGA Journal*, *36*(2), 67–80.

Frömer, R., Lin, H., Dean Wolf, C., Inzlicht, M., & Shenhav, A. (2021). Expectations of reward and efficacy guide cognitive control allocation. *Nature Communications*, *12*(1), 1030.

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278.

Gobet, F. (1997). A pattern-recognition theory of search in expert problem solving. *Thinking & Reasoning*, *3*(4), 291–313.

Gobet, F., & Jansen, P. (1994). Towards a chess program based on a model of human memory. *Advances in Computer Chess*, *7*, 35–60.

Grahek, I., Frömer, R., Prater Fahey, M., & Shenhav, A. (2023). Learning when effort matters: Neural dynamics underlying updating and adaptation to changes in performance efficacy. *Cerebral Cortex*, *33*(5), 2395–2411.

Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, *7*(2), 217–229.

Hall-McMaster, S., Muhle-Karbe, P. S., Myers, N. E., & Stokes, M. G. (2019). Reward boosts neural coding of task rules to optimize cognitive flexibility. *Journal of Neuroscience*, *39*(43), 8549–8561.

Holding, D. H. (1992). Theories of chess skill. *Psychological Research*, *54*(1), 10–16.

Horvitz, E. J. (1991). *Computation and action under bounded resources*. Stanford University.

Huys, Q. J., Eshel, N., O'Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: how the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS computational biology*, *8*(3), e1002410.

Kool, W., Gershman, S. J., & Cushman, F. A. (2017). Cost-benefit arbitration between multiple reinforcement-learning systems. *Psychological Science*, *28*(9), 1321–1333.

Kuperwajs, I., Ho, M. K., & Ma, W. J. (2024). Heuristics for meta-planning from a normative model of information search. *PsyArXiv*.

Kuperwajs, I., van Opheusden, B., Russek, E. M., & Griffiths, T. L. (2024). Learning from rewards and social information in naturalistic strategic behavior. *PsyArXiv*.

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.

Mastrogiuseppe, C., & Moreno-Bote, R. (2022). Deep imagination is a close to optimal policy for planning in large decision trees under limited resources. *Scientific Reports*, *12*(1), 10411.

Mattar, M. G., & Lengyel, M. (2022). Planning in the brain. *Neuron*, *110*(6), 914–934.

Miller, K. J., & Venditto, S. J. C. (2021). Multi-step planning in the brain. *Current Opinion in Behavioral Sciences*, *38*, 29–39.

Pearl, J. (1984). *Heuristics: intelligent search strategies for computer problem solving*. Addison-Wesley Longman.

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. *PsyArXiv*.

Russell, S., & Wefald, E. (1991). Principles of metareasoning. *Artificial intelligence*, *49*(1-3), 361–395.

Sezener, C. E., Dezfouli, A., & Keramati, M. (2019). Optimizing the depth and the direction of prospective planning using information values. *PLoS computational biology*, *15*(3), e1006827.

Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., … others (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, *362*(6419), 1140–1144.

Sunde, U., Zegners, D., & Strittmatter, A. (2022). Speed, quality, and the optimal timing of complex decisions: Field evidence. *arXiv preprint arXiv:2201.10808*.

van Opheusden, B., Kuperwajs, I., Galbiati, G., Bnaya, Z., Li, Y., & Ma, W. J. (2023). Expertise increases planning depth in human gameplay. *Nature*, *618*(7967), 1000–1005.

van Opheusden, B., & Ma, W. J. (2019). Tasks for aligning human and machine planning. *Current Opinion in Behavioral Sciences*, *29*, 127–133.