

Speed, Quality, and the Optimal Timing of Complex Decisions: Field Evidence

Uwe Sunde, Dainis Zegners, Anthony Strittmatter

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Speed, Quality, and the Optimal Timing of Complex Decisions: Field Evidence

Abstract

This paper presents an empirical investigation of the relation between decision speed and decision quality for a real-world setting of cognitively-demanding decisions in which the timing of decisions is endogenous: professional chess. Move-by-move data provide exceptionally detailed and precise information about decision times and decision quality, based on a comparison of actual decisions to a computational benchmark of best moves constructed using the artificial intelligence of a chess engine. The results reveal that faster decisions are associated with better performance. The findings are consistent with the predictions of procedural decision models like drift-diffusion-models in which decision makers sequentially acquire information about decision alternatives with uncertain valuations.

JEL-Codes: D010, D900, C700, C800.

Keywords: response times, speed-performance profile, drift-diffusion model, uncertain evaluations.

Uwe Sunde
LMU Munich / Germany
uwe.sunde@lmu.de

Dainis Zegers
School of Management, Erasmus University
Rotterdam / The Netherlands
zegers@rsm.nl

Anthony Strittmatter
CREST-ENSAE
Paris / France
anthony.strittmatter@ensae.fr

January 26, 2022

Uwe Sunde gratefully acknowledges financial support of the Deutsche Forschungsgemeinschaft (CRC TRR 190 Rationality and Competition, project number 280092119). Anthony Strittmatter gratefully acknowledges financial support of the French National Research Agency (LabEx Ecodec/ANR-11-LABX-0047).

1 Introduction

Many real-world decisions in economics and management involve both substantial complexity and time pressure. The question of whether slow decisions or fast decisions are associated with higher quality remains an open issue, however. The relation between the speed and the quality of decisions in settings in which the timing of choices is endogenous is unclear because different mechanisms work in opposite directions. Taking more time to reach a decision can help to make a more thoughtful choice, implying that slower decisions are of higher quality. Alternatively, slow decisions might entail lower quality if they reflect problems that decision makers perceive as more difficult, in terms of higher complexity or less clear information about the implications of different decision alternatives. Moreover, fast decisions can be of high quality when the decision maker has a strong intuition about the right choice or relatively precise information early on during the decision-making process so that the value of additional cogitation or information acquisition is relatively low. Existing evidence on the relation between the endogenous timing of decisions and their quality is scarce and largely limited to stylized experiments that often relate to fairly simple choices.

This paper presents an empirical investigation of the relation between decision speed and decision quality for cognitively-demanding real-world decisions in a competitive environment with high stakes: professional chess. Conceptually, chess is ideally suited to investigate this relation due to the extraordinarily rich and precise information about the context in which decisions are made. In particular, in this setting the complexity and difficulty of decisions exhibits considerable variation that can be measured with high accuracy, and the timing of decisions is endogenous. The decision making process in chess can be viewed as the solution to a complex problem in which the decision maker is uncertain about the quality of a particular move and, due to time restrictions, faces costs for gathering more information by cogitating over the various decision alternatives to determine the optimal move. The choice problem can thus be characterized as the determination of a move by an optimal stopping rule in a setting with uncertain priors about the quality of alternative moves. Conceptually, this closely resembles the decision

problems often faced by managers in practice. Professional chess thus represents a unique source of high-resolution data from a non-experimental setting about behavior of decision makers facing high stakes, who are observed repeatedly while solving complex choice problems under time pressure, and with varying difficulty.

The empirical analysis is based on information for decision times and move quality for more than 80,000 moves from almost 1,600 games between professional chess players. For each move, the configuration of pieces on the board is analyzed with a state-of-the-art chess engine, which provides a precise and objective estimate of the best possible move for a player facing this configuration. For each move, the speed of the decision made by the player is available at high resolution. The use of artificial intelligence built into chess engines allows constructing clean and objective measures for the quality of a particular decision and for other relevant aspects such as the difficulty of a given decision problem. This makes it possible to measure decision time and decision quality with high accuracy. Moreover, the use of move-by-move data implies that the analysis is able to reveal general patterns of decision-making while absorbing systematic heterogeneity across individual decision makers.

The estimation results document two sets of findings. First, regarding the determinants of decision speed, the results show that decisions are faster when the remaining time budget of a player is more constrained, implying greater time pressure and hence a greater marginal cost associated with additional cogitation. Likewise, decisions are slower if the problem is objectively more complex or if the uncertainty about the best decision is greater. Second, faster decisions are robustly associated with higher decision quality. In addition, and conditional on decision times, lower decision quality is associated with higher time pressure, greater complexity and greater uncertainty.

Contribution to the Literature. The first set of findings corresponds to the intuitive predictions of models of decision making that decision makers spend more time, which is a proxy for the effort devoted to cogitation, when the benefits of cogitation exceed its costs (see, e.g., Gabaix et al., 2006; Alaoui and Penta, 2016). By considering field

data with within-subject variation in decision times conditional on fine-grained variation in time pressure, complexity, and uncertainty, our findings complement earlier lab findings of longer consideration times when the gap in the subjective evaluation between two options is relatively small (Chabris et al., 2009), findings pointing at the sensitivity of reaction time with respect to strength of preferences (see, e.g., Krajbich et al., 2015; Konovalov and Krajbich, 2019), recent work on deliberation times as correlate of cognitive effort influencing the depth and the intensity of reasoning (Alós-Ferrer and Buckenmaier, 2021), and work on adaptive decision making in strategic environments (Spiliopoulos et al., 2018).

Our second set of findings contributes to the small literature on the relation between endogenous decision times and decision performance, and on the underlying processes of reasoning in strategic decision contexts (e.g., Spiliopoulos and Ortmann, 2018, for a recent survey). Response times have been studied based on notions of two-system models distinguishing between instinctive and cognitive (or contemplative) decision-making (Arad and Rubinstein, 2012; Spiliopoulos et al., 2018), and for classifying individuals into instinctive and contemplative types (Rubinstein, 2007, 2016). While this work typically focuses on observations of behavior of one individual in different games, our data comprise many observations for decisions of the same individual in the same decision context, but with varying complexity and time pressure. Only few studies have used repeated strategic interactions in the lab to measure the strategic complexity of a decision, with findings suggesting that decision times are responsive to strategic complexity and suggest that “overthinking” (i.e., thinking longer than usual) leads to lower decision quality when relating decision quality to decision times relative to decision times in similar decision settings (see, e.g., Gill and Prowse, 2019, for a recent contribution using a p -beauty contest with feedback). Our finding of a negative relation between decision time and decision quality is overall consistent with this, but our results are based on within-subject comparisons using variation in the complexity of decisions and in time pressure. By explicitly accounting for variation in difficulty and time pressure, our analysis addresses concerns about these as potential confounders in the respective literature

(see, e.g., Krajbich et al., 2015; Spiliopoulos and Ortmann, 2018, for a discussion of this point and the need for incorporating difficulty).

Our findings are also informative about the underlying process of decision making. In particular, the negative relation between the time spent on the decision and the quality of the decision is inconsistent with simple two-system models of decision making that predict that faster decisions due to, e.g., greater time pressure and lower subjective importance, are less likely to be based on deep reasoning and thus exhibit lower quality in complex settings (see, e.g., Kahneman, 2003; Dane and Pratt, 2007). Instead, the negative correlation is consistent with two-system models in which decision speed and quality depend on the alignment of the solutions based on an intuitive system of decision-making (that might be based on experience accumulated in the past) and cogitation (Achtziger and Alós-Ferrer, 2014; Sahm and von Weizsäcker, 2016; Alós-Ferrer, 2018; Spiliopoulos, 2018). Earlier work has predicted that a greater time budget leads to better decisions by allowing for processing more decision alternatives (see, e.g., Busemeyer and Townsend, 1993). Our findings are consistent with this prediction and confirm findings that additional time available for deliberation improves performance (see, e.g., Chabris and Hearst, 2003; Moxley et al., 2012). However, our results also indicate the need to account for decision times when analyzing the role of time pressure on decision quality, since the positive relation between remaining time and decision quality only holds once conditioning on decision time. This is consistent with recent findings that the observed distribution of choices and their correlates critically depends on controlling for response time (Webb, 2019). Our results allow disentangling the roles of available time budget (which reflects the implicit costs for later decisions), and the complexity and uncertainty associated with the decision, for the speed and quality of the decision.

Our empirical results are also consistent with the predictions of recent work on drift-diffusion models that has considered settings in which the relative evaluations of decision alternatives are uncertain ex-ante and that has shown that the positive association between decision speed and decision quality dominates in such settings (see Fudenberg et al., 2018, 2020, for details). So far, drift-diffusion models have been mainly used

to model decisions in simple decision problems and existing evidence is restricted to lab experiments with non-complex decisions (see, e.g., Ratcliff and Rouder, 1998; Krajbich and Rangel, 2011; Ratcliff et al., 2015; Bhui, 2019; Schotter and Trevino, 2021; Huseynov and Palma, 2021). Our results document that the relevance of drift-diffusion models with uncertain evaluations of decision alternatives extends to complex decision contexts. Moreover, our findings provide evidence consistent with the underlying mechanisms. To our knowledge, the present study is the first to provide evidence from a realistic field setting with experienced decision makers and an objective measure of decision quality. On a more applied side, our results inform the literature on intuitive decision making by experts (see, e.g., Salas et al., 2010, for a review). In this context, our results complement empirical work on bounded rationality (Zegners et al., 2020), suggesting that human intuition and experience are important factors in determining optimal decisions.

Our analysis continues with some conceptual considerations and a description of the data and the empirical measures in Section 2. Section 3 contains the results. In Section 4, we provide a discussion of the results and show that the negative relation between time spent on a decision and the quality of the decision is also consistent with the predictions of drift-diffusion-models.

2 Empirical Implementation

2.1 Conceptual Considerations

The analysis of the relationship between decision speed and decision quality is rooted in a conceptual setting in which a decision maker is confronted with several decision alternatives and faces time pressure. Concretely, the decisions in the empirical setting of chess involve making a decision about the next move in a particular configuration of pieces on the board under a limited time budget. Decisions involve the choice of a particular move out of a set of alternative moves, which corresponds to a discrete choice problem. The underlying neural dynamics of decision making imply that decision speed

and decision quality are closely related. An approach to model these decision dynamics used in bounded accumulation models from psychology is to assume that the cogitation process corresponds to the sequential acquisition and processing of information about the consequences of different move alternatives, and an optimal stopping rule. This implies that the observed distribution of decisions and their respective quality crucially depends on decision times.

Empirical assessments of the relation between decision times and decision quality typically focus on particular regularities such as individual differences in response time and accuracy, often using between-subject designs or within-subject designs with only few choices. This raises identification problems related to unobserved heterogeneity across decision makers and decision problems. The ideal setting to investigate the relationship between decision quality and decision time involves a within-subject design with many observations of similar, but not identical, decisions. The chess data used in the empirical analysis below provide such a setting of within-subject variation with observations of a large number of decisions and the associated decision speed and quality for the same decision maker. By their nature, moves in chess constitute full information rational choice problems under a time constraint and with uncertainty about the best possible solution. Moreover, in chess, decision makers only rarely face the exact same choice problem, especially when abstracting from the opening phase of the game, due to the abundance of different configurations of pieces on the board, which creates considerable variation in time pressure and in the difficulty of the decision. Hence, the empirical implementation is based on information about the decisions of the same individual across choice problems that differ due to the difference in the configuration on the board.

Our analysis proceeds in two steps. The first step refers to the factors that influence decision times, and the second step considers the relation between the quality of decisions and decision times as well as the other decision-relevant factors. In the first step, we consider time pressure and the difficulty of the decision as main determinants of decision times. The testable hypotheses are evaluated using proxies for these determinants. In particular, as a proxy for time pressure, we use the remaining time budget

for a given player in a given configuration of the game. A smaller time budget implies a larger opportunity cost since spending an additional unit of time on cogitating over a move implies less time available for future moves. This constraint becomes increasingly binding as the remaining time budget converges to zero. Hence, greater time pressure (in terms of a smaller remaining time budget) implies higher costs for later decisions in the game and should therefore lead to faster decisions. Coming up with proxies for the difficulty of the decision problem is more complicated since difficulty can arise for distinct reasons. In particular, greater difficulty can arise from greater positional complexity, where more complex problems imply a slower process until a decision can be reached with sufficient confidence. This can be either due to a slower process of cogitation and information acquisition, or due to a more noisy and unpredictable cogitation process. As a proxy for this feature, we use a measure of complexity of a given configuration that is based on the number of alternative moves that need to be considered in a given positional setting. Alternatively, difficulty can arise from greater uncertainty related to the evaluation of alternatives. This affects the possibility to discriminate between alternatives and hence the uncertainty associated with a decision. To capture this dimension, we use information about the evaluation gap between the best conceivable move for a given configuration, and the second-best move. Greater difficulty, either in terms of greater complexity of the decision or in terms of a more diffuse prior about the optimal decision, is expected to lead to longer decision times until a choice is made.

In the second step, we investigate the relation between decision quality and decision times, as well as the role of the factors that influence decision times. The relation between decision quality and decision times in complex, cognitively-demanding tasks is an empirical question and depending on the particular model, different predictions are possible. A similar statement applies to the influence of time pressure and the difficulty of the decision. It is likely, however, that the influence of these factors on decision quality crucially depends on the timing of decisions.

2.2 Data and Measurement

Data from Professional Chess Players. The data used in the empirical analysis have been collected from an internet platform that broadcasts all professional over-the-board chess tournaments (www.chess24.com) and contains detailed information for more than 80,000 moves from around 1,600 games that were played in 54 single round-robin tournaments during the years 2014-2017. These include events such as national championships, qualifications tournaments for the world championship and the most prestigious invitational tournaments in the chess calendar. Table A1 in the Appendix provides a list of tournaments included in the data set. All games were played at regular time controls that allocate a time budget of a minimum of 2 hours thinking time to each player to conclude the game.¹ The data contain detailed information about the players, including their performance statistics in terms of their ELO number.² We restrict our analyses to games between professional players with an ELO number of at least 2,500 at the time of the game.³

The data include precise information about the time consumed for each move and about the remaining time budget. In addition, the move-by-move data comprise information about the exact configuration of pieces on the board. We use this information to compute an evaluation of this configuration in terms of the relative standing of each player, an evaluation of the complexity of the configuration, and an evaluation of move quality, using a chess engine as explained in more detail below. We exclude the first fifteen moves of each player in a game from all our analyses. These moves are typically

¹According to the regulations by the International Chess Federation FIDE, for a game to be rated for determining the world rankings each player must have a minimum of 120 minutes, assuming the game lasts 60 moves per player (see <https://handbook.fide.com>, last accessed January 3, 2022). Tournaments that are not officially organized by FIDE use slight variations of the official FIDE time control regime.

²The ELO number constitutes a method for calculating the relative playing strength of players (invented by the Hungarian mathematician Arpad Elo). According to international conventions, an ELO number of at least 2,500 is a requirement for being awarded the title of an international grandmaster (this requirement has to be fulfilled once during the career, but does not have to be maintained to keep the title, see <https://handbook.fide.com/chapter/B01Regulations2017>, last accessed January 3, 2022).

³We also excluded games with decision times that are likely to be mismeasured (such as negative decision times or times of several hours devoted to a single move) and all games from tournaments that included multiple games exhibiting problematic decision time data.

the result of routine openings, which are studied intensively and memorized by players in the preparation for a game.

Measuring Decision Quality. To construct a measure of decision quality, we make use of a chess engine that computes the best possible move for a given configuration of pieces on the chessboard. In particular, we use the open-source chess engine STOCKFISH, which is considered to be one of the best available programs and has an estimated ELO rating of approximately 3544 points.⁴ In comparison, the incumbent World Champion Magnus Carlsen had an ELO rating of 2856 points in December 2021 according to the official rating list by the International Chess Federation FIDE.

For each configuration, the chess engine creates a game-tree for all possible moves by white and black for a pre-specified depth of n moves ahead, the so-called search depth. We restrict the engine to a search depth of 24 moves to economize on computing costs. Then, the configurations at the respective end-nodes are evaluated in terms of pieces left on the board, safety of the king, mobility of pieces, pawn-structure, etc. Based on this evaluation, the engine then determines the best move using an algorithm conceptually similar to backward induction under the assumption of mutually best responses.⁵

Concretely, the chess engine delivers a measure of the relative standing for a given

⁴We use the version STOCKFISH 14 with a strength of play that is equivalent to ELO of 3544 points according to <https://ccr1.chessdom.com/ccr1/4040/> (accessed on December 20, 2021).

⁵Most modern chess engines are based on domain-specific algorithmic heuristics that were developed specifically to search the sequential game-tree arising from a given configuration. In particular, these chess engines use an enhanced version of the min-max algorithm that disregards branches of the search tree that have already been found to be dominated. This reduces the search-space without impacting the final choice of the best move by the engine (<https://www.chessprogramming.org/Alpha-Beta>, last visited April 18, 2021). Only recently, more general machine learning techniques in the form of neural networks have been embodied in chess engines such as Google’s non-public engine *AlphaZero* (Silver et al., 2018) and the open-source engine *LeelaChessZero*. We use a version of *Stockfish 14* that searches the game tree using the standard approach of Alpha-Beta search and evaluates the end positions at the end nodes using a neural network trained on a large database of pre-evaluated chess positions (<https://stockfishchess.org/blog/2020/introducing-nnue-evaluation>, last visited January 5, 2022). It is very likely that *Stockfish 14* is considerably stronger than *AlphaZero* based on the fact that Stockfish’s current ELO is estimated to be by about 200 points higher than the version of Stockfish that was defeated by *AlphaZero* in 2018 (see <https://chess.stackexchange.com/questions/29791/which-is-better-stockfish-10-or-alphazero> (accessed January 17, 2022)). Since Google’s *AlphaZero* is non-public and has not played any publicly released games since 2018, this assessment cannot be validated, however.

configuration of pieces on the board. This measure reflects an evaluation of the current position of a player and represents a proxy of the winning odds. The evaluation of the current position is the result of the engine computing, for each configuration observed in the data set, the best continuation, and the relative standing is measured in terms of so-called pawn units, where one unit approximates a continuation that is associated with an advantage of possessing one more pawn.⁶ Based on this information, we compute a measure of performance, in terms of the quality of play of a given player in a given positional setting. This measure is based on the comparison of the actual move made by the respective player to the best conceivable move as suggested by the chess engine. This move is not necessarily the absolute best move, but on average the move suggested by the engine is better than conceivable by any human player. Decision quality is measured by a binary indicator of whether a player makes the best possible move as suggested by the chess engine in a given configuration. In robustness checks, we also computed the deviation of a player’s move (in terms of pawn units) from the best move identified by the chess engine.

Proxy Variables of Decision-relevant Factors. We use the available information to construct three proxy measures for factors that are central to the empirical analysis of the testable hypotheses. First, we use the time budget in terms of remaining time as a proxy of the (inverse) time cost for cogitation. Second, we use a chess engine to compute, for each observed configuration, a measure of complexity of the configuration. This measure corresponds to a proxy for the difficulty of the decision. The more complex the configuration, the more nodes the chess engine needs to search in the game-tree arising from a given position. The number of nodes needed for the engine to compute the best strategy for the next n moves ahead therefore serves as a measure of the complexity of a given configuration.⁷ In particular, we use the number of *mega-nodes* (number of

⁶This measure is relative. Positive numbers indicate an advantage for the player with white pieces, negative numbers indicate an advantage for the player with black pieces. For example, if the evaluation is -1.00 pawn units, black is better “as if one pawn up.”

⁷As a baseline measure of complexity, we use the number of nodes computed by the chess engine to reach a search-depth of 24 moves.

nodes divided by 1,000,000) as a measure of complexity.⁸ Third, we compute a proxy of the (inverse of) uncertainty of the prior, in terms of the difference between the best continuation as stipulated by the chess engine and the second-best continuation. The greater this difference, the smaller the uncertainty about the relative evaluation of the best move and any other (i.e., the second-best) alternative; the smaller this difference, the more diffuse the prior (and the greater the uncertainty) about the best continuation move.

Appendix Table A2 documents the descriptive statistics of the move-by-move data used in the analysis.

3 Empirical Results

Decision Speed. The first set of results relates to the determinants of decision time. Table 1 presents the results of multivariate regressions that include controls for (interacted) player-game fixed effects, the evaluation of the current position from the perspective of the decision maker, and for previously played moves (fixed effects for the number of moves played) to capture systematic heterogeneity that is not related to the research question.

The first finding, presented in Column (1), shows that decision makers spend more time on a move when the remaining time budget is larger. The results in Column (2) show that a greater complexity of a given configuration, as proxied by the number of mega-nodes computed by the chess engine to determine the best move for a given search depth, is associated with longer decision times. The results in Column (3) document that decision times are longer if the evaluation gap (i.e., the difference in the evaluation) of the best and of the second-best move is smaller. A smaller evaluation gap serves as proxy about greater uncertainty regarding the quality of alternative moves. These results also emerge when estimating a specification that includes all covariates, as shown in Column (4). The corresponding results for non-linear relationships between decision time and

⁸Modern engines like STOCKFISH calculate approximately 10-100 million nodes per second on standard personal computing hardware.

Table 1: Determinants of Decision Time

	Dependent Variable: Time Spent on Move (minutes)			
	(1)	(2)	(3)	(4)
<i>Time Budget</i>				
Remaining time (min.)	0.0985*** (0.0032)			0.0947*** (0.0032)
<i>Complexity</i>				
N Mega-nodes computed		0.0336*** (0.0031)		0.0193*** (0.0028)
<i>Evaluation Gap</i>				
Distance second best move (log)			-0.9195*** (0.0241)	-0.8294*** (0.0233)
Player-Game Fixed Effects	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Number Move Fixed Effects	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Control Evaluation Position	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Move Observations	80359	80359	80359	80359
Game Observations	1592	1592	1592	1592

Note: OLS estimates. The variable *Distance second best move (log)* is computed as $\ln(d+1)$ where d is the absolute difference between the evaluation of the best and the second best move in terms of pawn units as given by the chess engine. The evaluation of the current position is controlled for using two dummy variables indicating whether the current position is evaluated as better (> 0.5 pawn units) or worse (< -0.5 pawn units) for the player to move. Standard errors are clustered on the game level. *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$.

the respective measures are shown in Figure 1. Considering more flexible specifications confirms the main findings: Greater time pressure, less complexity and less uncertainty all imply faster decisions. These findings could be due to higher costs for cogitating, or due a faster cogitation process that induce faster decisions, respectively.

These results are robust to using alternative measures for the proxy variables for time pressure, complexity and the uncertainty regarding the best move (see Table A3 in the Appendix).⁹ The results are also robust to the estimation of an extended specification that includes controls for decision times and decision quality in previous moves to control

⁹In particular, as alternative proxy for time pressure we use an indicator for immediate time pressure characterizing moves before the time control at 40 moves. As an alternative proxy for complexity, we use the number of legal moves that can be executed in a given position. As an alternative proxy for uncertainty associated with move alternatives we use the standard deviation of the evaluation of the six best continuation moves suggested by the chess engine in a given position.

for a potential correlation within games (see Table A4 in the Appendix).




Figure 1: Determinants of Decision Time - Flexible Specifications

Note: OLS results of flexible specifications. Dependent variable is Time Spent on Move (minutes). All three graphs are based on estimates of flexible dummy specifications of variables from the same regression model that includes the same controls and fixed effects as the specifications in Table 1. Variables depicted on the horizontal axis are split into equally spaced intervals. Dots report point estimates and whiskers report 95% confidence intervals. Panel (a): Remaining Time measured in minutes; omitted reference category is Remaining Time < 5 minutes. Panel (b): Complexity of chess position as measured by number of mega-nodes (1 mega-node equals 1,000,000 nodes) computed by chess engine to reach search depth of 24 moves ahead; omitted reference category: < 1 mega-nodes computed. Panel (c): Distance second best move is computed as $\ln(d+1)$ where d is the absolute difference between the evaluation of the best and the second best move in terms of pawn units as given by the chess engine; omitted reference category is distance second best move < 0.1.

Decision Quality. The second set of results refers to the quality of decisions and its association with decision speed. Table 2 presents the results from multivariate regressions with move quality, as proxied by a binary indicator of whether the best move is played, as dependent variable. All specifications include controls for (interacted) player-game fixed effects, current positions, and fixed effects for the number of previously played moves. Column (1) presents the empirical results for the association between decision quality and speed. The findings show that faster decisions (when less time is spent on the decision) are associated with greater decision quality (a higher propensity of making the best possible move). The results in Columns (2)-(4) refer to the relation between decision quality and the determinants of decision speed in isolation. For each of the determinants, the sign of the respective coefficient corresponds to a negative association between the implied time for a decision and decision quality, consistent with the results of Table 1.

Column (5) presents the results for a specification with decision time and the different proxies. The results reveal a more subtle picture. In this specification, the negative association between decision time and decision quality is confirmed. In addition, the results reveal independent effects of the time budget, complexity and the evaluation gap on decision quality. In particular, once conditioning on decision time, the sign of the effect of remaining time budget turns positive and significant. This is consistent with the hypothesis that, conditional on the decision time, a larger time budget is associated with higher decision quality. This suggests an independent effect of lower cogitation costs, potentially due to less emotional stress or cognitive load. This effect is unlikely to be due to physiological factors related to, e.g., better concentration when fatigue is lower, because such factors are captured by the control for the number of previously played moves. Conditional on decision time, greater complexity is associated with a lower quality of the decision on average. Likewise, greater uncertainty (as proxied by a smaller evaluation distance) is associated with a lower decision quality. The distinct influence of response time, remaining time, and difficulty (in terms of complexity and the evaluation gap) provide support for concerns about difficulty as a potential confounder

Table 2: Determinants of Decision Quality


	Dependent Variable: Best Move (Dummy)					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Decision Time</i>						
Time spent on move (min.)	-0.0227*** (0.0006)				-0.0164*** (0.0005)	-0.0189*** (0.0010)
<i>Time Budget</i>						
Remaining time (min.)		-0.0019*** (0.0002)			0.0006*** (0.0002)	0.0007*** (0.0002)
<i>Complexity</i>						
N Mega-nodes computed			-0.0088*** (0.0005)		-0.0056*** (0.0004)	-0.0050*** (0.0004)
<i>Evaluation Gap</i>						
Distance second best move (log)				0.2132*** (0.0036)	0.1931*** (0.0034)	0.1839*** (0.0035)
<i>Interactions</i>						
Time spent × N Mega-nodes computed						-0.0002** (0.0001)
Time spent × Dist. second best move						0.0134*** (0.0020)
Player-Game Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Number Move Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Control Evaluation Position	Yes	Yes	Yes	Yes	Yes	Yes
Move Observations	80359	80359	80359	80359	80359	80359
Game Observations	1592	1592	1592	1592	1592	1592

Note: OLS estimates. The variable *Distance second best move* is computed as $\ln(d + 1)$ where d is the absolute difference between the evaluation of the best and the second best move in terms of pawn units as given by the chess engine. The evaluation of the current position is controlled for using two dummy variables indicating whether the current position is evaluated as better (> 0.5 pawn units) or worse (< -0.5 pawn units) for the player to move. Standard errors are clustered on the game level. *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$.

in the relation between response times and performance (see Krajbich et al., 2015; Spiliopoulos and Ortmann, 2018).


Figure 2 presents corresponding non-linear relationships between decision quality and decision speed, as well as the proxy variables for the model parameters.

Column (6) of Table 2 presents results for an extended specification that includes interaction terms between time spent and the proxies for the difficulty of a move (in terms of complexity and the uncertainty that is associated with a smaller evaluation



(a) Decision Speed: Time Spent on Move

(b) Time Budget: Remaining Time



(c) Complexity: Seconds to reach fixed search depth

(d) Evaluation Gap: Distance second best move

Figure 2: Determinants of Decision Quality - Flexible Specifications

Note: OLS results of more flexible specifications. Dependent variable is Best Move (dummy). All graphs are based on estimates of flexible dummy-specifications of variables from the same regression model that includes same controls and fixed effects as the specifications in Table 2. Variables depicted on the horizontal axis are split into equally spaced intervals. Dots report point estimates and whiskers report 95% confidence intervals. Panel (a): Decision Speed measured in time spent on move in minutes; omitted reference category: time spend < 0.5 minutes. Panel (b): Remaining Time measured in minutes, reference category is Remaining Time < 5 minutes. Panel (c): omitted reference category: Mega-nodes computed to reach fixed search depth < 1 . Panel (d): Distance second best move is computed as $\ln(d + 1)$ where d is the absolute difference between the evaluation of the best and the second best move in terms of pawn units as given by the chess engine; omitted reference category is Distance second best move < 0.1 .

gap). While the main effects remain qualitatively and quantitatively almost unchanged, the interaction terms are significant. This suggests that the effects of complexity and uncertainty interact with decision time in determining the quality of decisions. In par-

ticular, the negative effect of longer decision times on decision quality is amplified by greater complexity, implying that longer cogitation does not necessarily attenuate the detrimental effect of complexity on decision quality. Similarly, the negative effect of longer decision times is aggravated the greater the uncertainty about the best move (in terms of a smaller distance in the evaluation of the best and second-best moves). This suggests that longer cogitation does not help in finding the best move in complex settings or when priors about the best decision are diffuse. In such situations, intuition and experience seem to play an important role for the determination of the optimal decision.

The main results are robust to using alternative measures for the cost of deliberation, complexity and the evaluation gap (see Table A5 in the Appendix). They are also robust to including controls for measures of decision times and decision quality in previous moves to account for potential correlation within games (see Table A6 in the Appendix) and when using a continuous measure for move quality along the intensive margin of move quality in terms of pawn units (see Table A7 in the Appendix). The coefficients of the interaction terms of decision time with complexity and uncertainty reveal opposite signs when considering decision quality along the intensive margin as dependent variable, however (see Table A7 in Appendix).

4 Discussion

The evidence presented in the previous section presents two sets of results on the determinants of decision speed and the relation between decision speed and decision quality in a real-world setting of complex and cognitively demanding decisions, using high-frequency data from professional chess. The first set of results documents that greater time pressure leads to faster decisions, while more difficult decisions and decisions that are associated with greater a priori uncertainty about the best alternative lead to slower decisions. The second set of results documents a positive relationship between decision speed and decision quality, using an empirical strategy for decision quality that is based on a benchmark that utilizes the artificial intelligence embodied in chess engines.

These findings are consistent with the predictions of drift-diffusion models with uncertain evaluations of decision alternatives. While these models have been considered in the context of simple choices (like choosing between two different candy bars), our results document that the predictions extend to complex strategic decision contexts and provide evidence on the underlying mechanisms. In drift-diffusion models, a sub-class of bounded accumulation models, noisy information is accumulated until a response criterion (a “boundary”) is reached. Once this criterion, which reflects the amount of evidence needed to make a decision, is reached, the decision is made accordingly (see, e.g., Bogacz et al., 2006; Ratcliff et al., 2015). Recent work has shown that bounded accumulation models have an analogue representation as a random utility model of discrete choice (Webb, 2019). An important implication of this equivalence is that the observed distribution of decisions (or preferences) crucially depends on decision times.

In the present context, the decision regarding the optimal continuation move for a given configuration on the board corresponds to the situation considered in drift-diffusion models with uncertain evaluations. Information acquisition occurs in the form of cogitation, which is modelled as a diffusion process with drift and random stimuli, and a decision about a move is made when the cogitation process reaches decision a boundary.¹⁰ As consequence, the observation that a move is made is associated with a realization of decision time and of decision quality, which is measured as the probability of making the best possible move in a given configuration conditional on making the decision at a certain time. Uncertainty enters in terms of the evaluations of the choice alternatives. The decision maker has a prior belief about the evaluations, but is uncertain about the best move, either due to the analytical complexity of the problem or due to the strategic uncertainty. Hence, the decision maker incurs costs over time while observing (independ-

¹⁰Most existing treatments of the drift-diffusion model consider choices between two alternatives. The characterization below extends to all choice pairs in the set of feasible choices (see Fudenberg et al., 2020, for details). Recent work has generalized the model to multiple alternatives (see, e.g., Roxin, 2019) without affecting the main implications for the purposes of this study. Hence, for simplicity and concreteness, let the decision be represented either by the best possible move, or by any other move, i.e., the decision is modeled as binary choice $i \in X$ with $X = \{x, y\}$. The diffusion process is modeled as $Z_t = \delta t + \alpha B_t$ in continuous time t with drift δ and a stimulus process B_t that is a Brownian motion with a parameter α . With binary choices, there are two boundaries, b^x or b^y . With symmetric boundaries, x is chosen if $b \leq Z_t$ and y if $-b \geq Z_t$.

dent) signals about the true evaluation (which also have the form of drift plus random stimuli that follow a Brownian motion). Since the decision depends on the relative evaluation of the choice alternatives, the evaluation difference is a sufficient statistic for the decision. Hence, the larger the perceived difference in evaluations the more informative is the signal (the larger is the drift towards the boundaries). It has been shown that the decision boundaries vary with decision time in potentially non-monotone ways, but the decision quality decreases with decision time when aggregating many decisions of a decision maker who on average has correct priors about the difference in the evaluation of the respective choice alternatives (see Fudenberg et al., 2018, 2020, for details). In this setting, the optimal choice can be characterized as the maximization of the posterior belief about the best possible move, subject to the cost associated with cogitation per instant of time. Equivalently, the decision problem can be characterized as the minimization of the regret of a wrong choice and of the sampling cost.¹¹

The probability of making the best possible move and the associated decision time are then affected by three mechanisms. The first is related to cogitation and information acquisition and implies that stopping later to make a move is associated with more information, which tends to increase the probability of choosing the optimal move. At the same time, two selection effects work in the opposite direction. First, optimal decisions are associated with earlier endogenous stopping when signals are strong and thus more informative (i.e., they exhibit a larger drift and/or a less noisy stimulus process). Second, endogenous stopping occurs earlier in trials where subjects experienced a faster understanding of the optimal behavior (which is equivalent to a more discriminatory prior and thus more informative signals). This selection effect only emerges in drift-diffusion models in which the evaluations associated with choices (i.e., their quality) are uncertain.

The characterization of the optimal stopping time in such a setting implies that with a higher cogitation cost, decisions are made earlier, and that the probability of making the best move is non-increasing in cogitation cost conditional on decision time. This

¹¹See Proposition 2 of Fudenberg et al. (2018).

prediction follows from a slow-down of updating as the value of cogitation declines with the time spent on cogitating over a decision. It is consistent with the finding of greater remaining time (i.e., lower time pressure) being associated with slower decisions (see Table 1 and Figure 1(a)) and, conditional on decision time, with higher decision quality (see Table 2 (column 5) and Figure 2(b)). The decision time also depends on the relative strength of the drift process and the noisiness of the signal, which are proxied by complexity and the evaluation gap between move alternatives, respectively. More informative signals are predicted to imply faster drift and/or less noise, and thus, faster and better decisions. Again, this is in line with the findings for complexity and the evaluation gap for decision speed (see Table 1 and Figures 1(b) and (c), respectively) and decision quality (see Table 2 and Figure 2(c) and (d), respectively). Moreover, in this setting, the two selection effects dominate, implying that decision quality decreases with decision time. This prediction is consistent with the negative association between time spent on a move and the quality of the move (see Table 2 and Figure 2(a)). These theoretical predictions are robust to asymmetric priors, endogenously divided attention on the two decision alternatives, or can be extended to non-linear cogitation (in particular increasing) cost (see Fudenberg et al., 2018, for details). In sum, the empirical results regarding decision time and decision quality are consistent with the predictions of uncertain-difference drift diffusion models.

Additional findings suggest that spending more time on a decision aggravates the negative effect of greater complexity, while spending more time on a decision when the uncertainty about the best decision is lower improves decision quality. These findings might be an indication of additional behavioral aspects that determine decision making as reflected in more elaborate models, such as dual-process diffusion models that account for alignment or conflict in the inference about the optimal decisions from a heuristic model of decision making and a maximizing model of decision making (see, e.g., Alós-Ferrer, 2018). The analysis here focused on the overall patterns of the relation between speed and decision quality in the context of endogenous timing of decisions and abstracted from analyzing systematic heterogeneity across individuals or the im-

plications of strategic interactions. For instance, to the extent that cognitive processes involved with decision making are affected by physiological factors, age might play a crucial role for cognitive performance (see, e.g., Strittmatter et al., 2020). Similarly, experience might interfere systematically with the way complexity or uncertainty shape decisions. A natural next step in the research agenda is to apply the methodology developed here to investigate the influence of individual-specific factors for the behavioral mechanisms in more detail. Disentangling the role of strategic interactions in shaping behavior constitutes another interesting direction for future research.

References

- ACHTZIGER, A. AND C. ALÓS-FERRER (2014): “Fast or Rational? A Response-Times Study of Bayesian Updating,” *Management Science*, 60, 923–938.
- ALAOUI, L. AND A. PENTA (2016): “Endogenous Depth of Reasoning,” *Review of Economic Studies*, 83, 1297–1333.
- ALÓS-FERRER, C. (2018): “A Dual-Process Diffusion Model,” *Journal of Behavioral Decision Making*, 31, 203–218.
- ALÓS-FERRER, C. AND J. BUCKENMAIER (2021): “Cognitive sophistication and deliberation times,” *Experimental Economics*, forthcoming, 24, 558–592.
- ARAD, A. AND A. RUBINSTEIN (2012): “Multi-dimensional iterative reasoning in action: The case of the Colonel Blotto game,” *Journal of Economic Behavior & Organization*, 84, 571–585.
- BHUI, R. (2019): “Testing Optimal Timing in Value-Linked Decision Making,” *Computational Brain & Behavior*, 2, 85–94.
- BOGACZ, R., E. BROWN, J. MOEHLIS, P. HOLMES, AND J. D. COHEN (2006): “The Physics of Optimal Decision Making: A Formal Analysis of Models of Performance in Two-Alternative Forced-Choice Tasks,” *Psychological Review*, 113, 700–765.
- BUSEMEYER, J. R. AND J. T. TOWNSEND (1993): “Decision Field Theory: A Dynamic-Cognitive Approach to Decision Making in an Uncertain Environment,” *Psychological Review*, 100, 432–459.
- CHABRIS, C. F. AND E. S. HEARST (2003): “Visualization, pattern recognition, and forward search: Effects of playing speed and sight of the position on grandmaster chess errors,” *Cognitive Science*, 27, 637–648.
- CHABRIS, C. F., D. LAIBSON, C. L. MORRIS, J. P. SCHULDT, AND D. TAUBINSKY (2009): “The Allocation of Time in Decision-Making,” *Journal of the European Economic Association*, 7, 628–637.
- DANE, E. AND M. G. PRATT (2007): “Exploring Intuition and its Role in Managerial Decision Making,” *Academy of Management Review*, 32, 33–54.
- FUDENBERG, D., W. NEWEY, P. STRACK, AND T. STRZALECKI (2020): “Testing the drift-diffusion model,” *Proceedings of the National Academy of Sciences*, 117, 33141–33148.
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2018): “Speed, Accuracy, and the Optimal Timing of Choices,” *American Economic Review*, 108, 3651–3684.

- GABAIX, X., D. LAIBSON, G. MOLOCHE, AND S. WEINBERG (2006): “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review*, 96, 1043–1068.
- GILL, D. AND V. PROWSE (2019): “Strategic Complexity and the Value of Thinking,” *Warwick University, CAGE Working Paper*, 431/2019.
- HUSEYNOV, S. AND M. A. PALMA (2021): “Food decision-making under time pressure,” *Food Quality and Preference*, 88, 104072.
- KAHNEMAN, D. (2003): “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review*, 93, 1449–1475.
- KONOVALOV, A. AND I. KRAJBICH (2019): “Revealed Indifference: Using Response Times to Infer Preferences,” *Judgment and Decision Making*, 14, 381–394.
- KRAJBICH, I., B. BARTLING, T. HARE, AND E. FEHR (2015): “Rethinking fast and slow based on a critique of reaction-time reverse inference,” *Nature Communications*, 6, 1–9.
- KRAJBICH, I. AND A. RANGEL (2011): “Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions,” *Proceedings of the National Academy of Sciences*, 108, 13852–13857.
- MOXLEY, J. H., K. A. ERICSSON, N. CHARNNESS, AND R. T. KRAMPE (2012): “The role of intuition and deliberate thinking in experts’ superior tactical decision-making,” *Cognition*, 124, 72–78.
- RATCLIFF, R. AND J. N. ROUDER (1998): “Modeling Response Times for Two-Choice Decisions,” *Psychological Science*, 9, 347–356.
- RATCLIFF, R., P. L. SMITH, AND G. MCKOON (2015): “Modeling Regularities in Response Time and Accuracy Data with the Diffusion Model,” *Current Directions in Psychological Science*, 24, 458–470.
- ROXIN, A. (2019): “Drift–diffusion models for multiple-alternative forced-choice decision making,” *Journal of Mathematical Neuroscience*, 9, 5.
- RUBINSTEIN, A. (2007): “Instinctive and Cognitive Reasoning: A Study of Response Times,” *Economic Journal*, 117, 1243–1259.
- (2016): “A Typology of Players: Between Instinctive and Contemplative,” *Quarterly Journal of Economics*, 131, 859–890.
- SAHM, M. AND R. K. VON WEIZSÄCKER (2016): “Reason, Intuition, and Time,” *Managerial and Decision Economics*, 37, 195–207.
- SALAS, E., M. A. ROSEN, AND D. DIAZGRANADOS (2010): “Expertise-Based Intuition and Decision Making in Organizations,” *Journal of Management*, 36, 941–973.

- SCHOTTER, A. AND I. TREVINO (2021): “Is response time predictive of choice? An experimental study of threshold strategies,” *Experimental Economics*, 24, 87–117.
- SILVER, D., T. HUBERT, J. SCHRITTWIESER, I. ANTONOGLU, M. LAI, A. GUEZ, M. LANCTOT, L. SIFRE, D. KUMARAN, T. GRAEPEL, ET AL. (2018): “A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play,” *Science*, 362, 1140–1144.
- SPILIOPOULOS, L. (2018): “The determinants of response time in a repeated constant-sum game: A robust Bayesian hierarchical dual-process model,” *Cognition*, 712, 107–123.
- SPILIOPOULOS, L. AND A. ORTMANN (2018): “The BCD of response time analysis in experimental economics,” *Experimental Economics*, 21, 383–433.
- SPILIOPOULOS, L., A. ORTMANN, AND L. ZHANG (2018): “Complexity, attention, and choice in games under time constraints: A process analysis,” *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44, 1609.
- STRITTMATTER, A., U. SUNDE, AND D. ZEGNERS (2020): “Life Cycle Patterns of Cognitive Performance over the Long Run,” *Proceedings of the National Academy of Sciences*, 117, 27255–27261.
- WEBB, R. (2019): “The (Neural) Dynamics of Stochastic Choice,” *Management Science*, 65, 230–255.
- ZEGNERS, D., U. SUNDE, AND A. STRITTMATTER (2020): “On the Causes and Consequences of Deviations from Rational Behavior,” *CRC 190 Rationality & Competition Discussion Paper*, 246.

Appendix with Supplementary Material For Online Publication

Additional Tables

Table A1: Tournaments included in Dataset

Tournament	Number Moves
National Championships	
Armenian Championship (2016 - 2017)	3994
Belarusian Championship (2017)	413
Dutch Championship (2015)	1239
French Championship (2016)	503
Icelandic Chess Championship (2016)	102
Kazakhstan Championship (2016)	103
Paraguay Championship (2016)	44
Polish Championship (2016)	2273
Russian Championship (2014, 2016)	5394
Russian Women's Championship (2016)	22
Swedish Championship (2016)	256
Swiss Championship (2016)	60
Ukrainian Championship (2016)	1401
World Chess Federation Tournaments	
FIDE Grand Prix - Baku (2014)	2334
FIDE Grand Prix - Khanty-Mansiysk (2015)	1986
FIDE Grand Prix - Tashkent (2014)	2610
FIDE Grand Prix - Tbilisi (2015)	2295
FIDE Women Grand Prix - Monte Carlo (2015)	1454
FIDE Women Grand Prix - Theran (2016)	1019
FIDE Women Grand Prix - Batumi (2016)	447
Invitational Tournaments	
Barcelona Masters (2016)	492
Batavia Chess Tournament (2016)	42
Dortmund Sparkassen Chess Meeting (2015 - 2016)	3028
GRENKE Chess Classic (2014)	552
Karpov Poikovsky Tournament (2016)	1749
Khanty-Mansiysk Women's Grand Prix (2016)	84
Lake Sevan (2016)	2409
Lev Gutman Birthday Jubilee Tournament (2015)	78
London Chess Classic (2014 and 2016)	1971
Margaryan Memorial (2017)	165
Norway Chess (2014 - 2016)	4536
Poikovsky (2015)	1964
Shamkir Chess (2014 - 2015)	3715
Sigeman & Co (2014)	106
Sinquefield Cup (2016)	1866
Suleymanpasa GM Tournament (2016)	216
Tal Memorial (2016)	2295
Tata Steel Challengers (2016 - 2017)	5595
Tata Steel Masters (2014, 2016 - 2017)	9052
U.S. Junior Championship (2016)	112
UKA German Masters (2015)	98
US.Championships (2015 - 2016)	5925
Vidmar Memorial (2016)	2568
Sum	80359

Note: This table shows an overview of tournaments that are included in our dataset. The tournaments include all elite single round-robin tournaments played in 2014 - 2017. Games with likely erroneous decision times such as negative decision times or decision times of several hours for a single move were excluded. Further, the data is restricted to games where both players have at least an ELO number of 2500. The data are taken from chess24.com.

Table A2: Descriptive Statistics – Move Level

	Moves	Mean	Std. Dev.	Min	Max
<i>Decision Quality</i>					
Best move (dummy)	80,359	0.565	0.496	0	1
Move quality (log-modulus)	80,359	-0.207	0.406	-5.784	0
<i>Decision Time</i>					
Time spent on move (min.)	80,359	2.712	4.241	0	63.67
<i>Time Budget</i>					
Remaining time (min.)	80,359	32.59	26.49	0.0167	140.8
Time pressure phase before move 40	80,359	0.240	0.427	0	1
<i>Complexity</i>					
N Mega-nodes computed	80,359	6.517	7.610	0.000873	241.8
N moves possible	80,359	29.95	11.78	2	67
<i>Evaluation Gap</i>					
Distance second best move (log)	80,359	0.501	0.786	0	5.957
Standard deviation evaluation 6 best moves	80,359	5.235	23.40	0	267.0

Note: Descriptive statistics for the baseline sample. The variable *Move Quality* is calculated as $\text{sign}(d) \cdot \ln(|d| + 1)$, where d measures the difference in the evaluation of the move played relative to the best possible move in pawn units. The variable *N moves possible* represents the number of possible legal moves in a given configuration. The variable *Distance second best move* is computed as $\ln(d + 1)$ where d is the absolute difference between the evaluation of the best and the second best move in terms of pawn units as given by the chess engine. The variable *Standard deviation evaluation 6 best moves* is calculated as the standard deviation of the evaluation of the chess positions according to the chess engine after each of the six best moves in a chess position would have been played in a given position.

Table A3: Determinants of Decision Time - Robustness - Alternative Measures

	Dependent Variable: Time spent on move (minutes)			
	(1)	(2)	(3)	(4)
<i>Time Budget</i>				
Time pressure phase before move 40	-1.3256*** (0.0411)			-1.2581*** (0.0411)
<i>Complexity</i>				
N moves possible		0.0541*** (0.0021)		0.0443*** (0.0021)
<i>Evaluation Gap</i>				
Standard deviation evaluation 6 best moves			-0.0122*** (0.0006)	-0.0080*** (0.0005)
Player-Game Fixed Effects	Yes	Yes	Yes	Yes
Control Move Number	Yes	Yes	Yes	Yes
Control Evaluation Position	Yes	Yes	Yes	Yes
Move Observations	80359	80359	80359	80359
Game Observations	1592	1592	1592	1592

Note: OLS estimates. The variable *Time pressure phase before move 40* is an alternative proxy for the time cost of cogitation and represents a dummy variable indicating whether the move number is in the range between move 30 and move 39. At move 40, each player receives additional time onto his time budget, therefore the moves before 40 are typically characterized as the time-pressure phase in chess. Since this does not allow to include fixed effects for the move number at the same time, all specification include a linear control variable for the move number. The variable *N moves possible* is an alternative proxy for complexity and represents the number of possible legal moves in a given configuration. A higher number of possible moves is typically associated with higher complexity. The variable *Standard deviation evaluation 6 best moves* is an alternative proxy for uncertainty of the prior and represents a measure of the standard deviation of the evaluation of the chess positions according to the chess engine after each of the six best moves in a chess position would have been played in a given position. The evaluation of the current position is controlled for using two dummy variables indicating whether the current position is evaluated as better (> 0.5 pawn units) or worse (< -0.5 pawn units) for the player to move. Standard errors are clustered on the game level. *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$.

Table A4: Determinants of Decision Time - Robustness - Previous Own Moves and Opponent

	Dependent Variable: Time Spent on Move (minutes)				
	(1)	(2)	(3)	(4)	(5)
<i>Time Budget</i>					
Remaining time (min.)	0.0947*** (0.0032)	0.0951*** (0.0032)	0.0929*** (0.0031)	0.0929*** (0.0031)	0.0926*** (0.0031)
<i>Complexity</i>					
N Mega-nodes computed	0.0193*** (0.0028)	0.0193*** (0.0028)	0.0197*** (0.0026)	0.0199*** (0.0026)	0.0199*** (0.0026)
<i>Evaluation Gap</i>					
Distance second best move (log)	-0.8294*** (0.0233)	-0.8291*** (0.0232)	-0.8286*** (0.0228)	-0.8332*** (0.0228)	-0.8296*** (0.0227)
<i>Decision Time Previous Moves</i>					
Time spent on move (previous own move)		0.0053 (0.0049)	-0.0125** (0.0049)	-0.0098** (0.0050)	-0.0087* (0.0050)
Time spent on move (previous opponent's move)			0.1313*** (0.0074)	0.1321*** (0.0074)	0.1264*** (0.0075)
<i>Decision Quality Previous Moves</i>					
Best move (previous own move)				0.1217*** (0.0292)	0.1289*** (0.0291)
Best move (previous opponent's move)					-0.2496*** (0.0309)
Player-Game Fixed Effects	Yes	Yes	Yes	Yes	Yes
Control Move Number	Yes	Yes	Yes	Yes	Yes
Control Evaluation Position	Yes	Yes	Yes	Yes	Yes
Move Observations	80359	80359	80359	80359	80359
Game Observations	1592	1592	1592	1592	1592

Note: OLS estimates. Column (1) shows for comparison the main specification from table 1. The evaluation of the current position is controlled for using two dummy variables indicating whether the current position is evaluated as better (> 0.5 pawn units) or worse (< -0.5 pawn units) for the player to move. Standard errors are clustered on the game level. *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$.

Table A5: Determinants of Decision Quality - Robustness - Alternative Measures

	Dependent Variable: Best Move (Dummy)					
	(1)	(2)	(3)	(4)	(5)	(6)
Time spent on move (min.)	-0.0215*** (0.0006)				-0.0207*** (0.0006)	-0.0222*** (0.0024)
<i>Time Budget</i>						
Time pressure phase before move 40		-0.0096* (0.0051)			-0.0432*** (0.0051)	-0.0430*** (0.0050)
<i>Complexity</i>						
N moves possible			-0.0057*** (0.0002)		-0.0036*** (0.0002)	-0.0037*** (0.0003)
<i>Evaluation Gap</i>						
Standard deviation evaluation 6 best moves				0.0023*** (0.0001)	0.0017*** (0.0001)	0.0017*** (0.0001)
<i>Interactions</i>						
N possible moves × Time spent						0.0000 (0.0001)
Standard deviation evaluation × Time spent						-0.0000 (0.0000)
Player-Game Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Control Move Number	Yes	Yes	Yes	Yes	Yes	Yes
Control Evaluation Position	Yes	Yes	Yes	Yes	Yes	Yes
Move Observations	80359	80359	80359	80359	80359	80359
Game Observations	1592	1592	1592	1592	1592	1592

Note: OLS estimates. The variable *Time pressure phase before move 40* is an alternative proxy for the time cost of cogitation, c , and represents a dummy variable indicating whether the move number is in the range between move 30 and move 39. At move 40, each player receives additional time onto his time budget, therefore the moves before 40 are typically characterized as the time-pressure phase in chess. Since this does not allow to include fixed effects for the move number at the same time, all specification include a linear control variable for the move number. The variable *N moves possible* is an alternative proxy for complexity and represents the number of possible moves in the current configuration. The variable *Standard deviation evaluation 6 best moves* is an alternative proxy for uncertainty of the prior and represents a measure of the standard deviation of the evaluation of the chess positions according to the chess engine after each of the six best moves in a chess position would have been played in a given position. The evaluation of the current position is controlled for using two dummy variables indicating whether the current position is evaluated as better (> 0.5 pawn units) or worse (< -0.5 pawn units) for the player to move. Standard errors are clustered on the game level. *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$.

Table A6: Determinants of Decision Quality - Robustness - Previous Own Moves and Opponent

	Dependent Variable: Best Move (Dummy)				
	(1)	(2)	(3)	(4)	(5)
<i>Decision Time</i>					
Time spent on move (min.)	-0.0164*** (0.0005)	-0.0164*** (0.0005)	-0.0169*** (0.0005)	-0.0170*** (0.0005)	-0.0169*** (0.0005)
<i>Complexity</i>					
N Mega-nodes computed	-0.0056*** (0.0004)	-0.0056*** (0.0004)	-0.0056*** (0.0004)	-0.0056*** (0.0004)	-0.0056*** (0.0004)
<i>Evaluation Gap</i>					
Distance second best move (log)	0.1931*** (0.0034)	0.1931*** (0.0034)	0.1926*** (0.0034)	0.1918*** (0.0034)	0.1917*** (0.0034)
Remaining time (min.)	0.0006*** (0.0002)	0.0005** (0.0002)	0.0005** (0.0002)	0.0005** (0.0002)	0.0005** (0.0002)
Time spent on move (previous own move)		-0.0013*** (0.0004)	-0.0018*** (0.0004)	-0.0013*** (0.0004)	-0.0014*** (0.0004)
Time spent on move (previous opponent's move)			0.0036*** (0.0004)	0.0037*** (0.0004)	0.0039*** (0.0004)
<i>Decision Quality Previous Moves</i>					
Best move (previous own move)				0.0207*** (0.0035)	0.0205*** (0.0035)
Best move (previous opponent's move)					0.0076** (0.0037)
Player-Game Fixed Effects	Yes	Yes	Yes	Yes	Yes
Number Move Fixed Effects	Yes	Yes	Yes	Yes	Yes
Control Evaluation Position	Yes	Yes	Yes	Yes	Yes
Move Observations	80359	80359	80359	80359	80359
Game Observations	1592	1592	1592	1592	1592

Note: OLS estimates. Column (1) shows for comparison the main specification without interactions from table 2. The evaluation of the current position is controlled for using two dummy variables indicating whether the current position is evaluated as better (> 0.5 pawn units) or worse (< -0.5 pawn units) for the player to move. Standard errors are clustered on the game level. *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$.

Table A7: Determinants of Decision Quality - Robustness - Intensive Margin

	Dependent Variable: Move Quality (Log-Modulus)					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Decision Time</i>						
Time spent on move (min.)	-0.0091*** (0.0004)				-0.0103*** (0.0004)	-0.0063*** (0.0008)
<i>Time Budget</i>						
Remaining time (min.)		-0.0001 (0.0002)			0.0009*** (0.0002)	0.0008*** (0.0002)
<i>Complexity</i>						
N Mega-nodes computed			-0.0082*** (0.0006)		-0.0084*** (0.0006)	-0.0089*** (0.0007)
<i>Evaluation Gap</i>						
Distance second best move (log)				-0.0259*** (0.0063)	-0.0429*** (0.0065)	-0.0307*** (0.0065)
<i>Interactions</i>						
Time spent × N Mega-nodes computed						0.0002** (0.0001)
Time spent × Dist. second best move						-0.0174*** (0.0028)
Player-Game Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Number Move Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Control Evaluation Position	Yes	Yes	Yes	Yes	Yes	Yes
Move Observations	80359	80359	80359	80359	80359	80359
Game Observations	1592	1592	1592	1592	1592	1592

Note: OLS estimates. The dependent variable Move Quality (Log-Modulus) is a measure of the distance of the actual move played from the optimal move stipulated by the chess engine, in terms of the log modulus transformation of pawn units, such that $Move\ Quality = sign(d) \cdot \ln(|d| + 1)$, where d measures the difference in the evaluation of the move played relative to the best possible move in pawn units. The evaluation of the current position is controlled for using two dummy variables indicating whether the current position is evaluated as better (> 0.5 pawn units) or worse (< -0.5 pawn units) for the player to move. Standard errors are clustered on the game level. *: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$.