



Response time distributions in rapid chess: a large-scale decision making experiment

Mariano Sigman^{1*}, Pablo Etchemendy¹, Diego Fernández Slezak² and Guillermo A. Cecchi³

¹ Physics Department, School of Sciences, University of Buenos Aires, Buenos Aires, Argentina

² Computer Science Department, School of Sciences, University of Buenos Aires, Buenos Aires, Argentina

³ Computational Biology Center, T.J. Watson Research Center, International Business Machines, Yorktown Heights, NY, USA

Edited by:

Shu-Chen Li, Max Planck Institute for Human Development, Germany

Reviewed by:

Victor H. de Lafuente, University of Washington School of Medicine, USA
Hugo Merchant, Universidad Nacional Autónoma de México, Mexico

*Correspondence:

Mariano Sigman, Physics Department, School of Sciences, University of Buenos Aires, Pabellón 1, Ciudad Universitaria, C1428EGA Buenos Aires, Argentina.
e-mail: sigman@df.uba.ar

Rapid chess provides an unparalleled laboratory to understand decision making in a natural environment. In a chess game, players choose consecutively around 40 moves in a finite time budget. The goodness of each choice can be determined quantitatively since current chess algorithms estimate precisely the value of a position. Web-based chess produces vast amounts of data, millions of decisions per day, incommensurable with traditional psychological experiments. We generated a database of response times (RTs) and position value in rapid chess games. We measured robust emergent statistical observables: (1) RT distributions are long-tailed and show qualitatively distinct forms at different stages of the game, (2) RT of successive moves are highly correlated both for intra- and inter-player moves. These findings have theoretical implications since they deny two basic assumptions of sequential decision making algorithms: RTs are not stationary and can not be generated by a state-function. Our results also have practical implications. First, we characterized the capacity of blunders and score fluctuations to predict a player strength, which is yet an open problem in chess softwares. Second, we show that the winning likelihood can be reliably estimated from a weighted combination of remaining times and position evaluation.

Keywords: games, machine learning, brain, intelligence, Markov, planning, brain-computer, parallel computing

INTRODUCTION

Chess has long been a model system to study complex thought processes (Groot, 1965; Charness, 1981; Holding and Reynolds, 1982; Gobet and Simon, 1996a; Schultetus and Charness, 1999; Reingold et al., 2001a,b). In particular, a consensus has emerged in that chess expertise comes in two forms: the ability to calculate variations (search) and the ability to recognize and remember meaningful patterns on the board (pattern recognition). The prevalent view is that expert players, as opposed to weaker ones, excel specifically at rapid object recognition abilities (Gobet and Simon, 1996a,b; Burns, 2004). Naively, one would expect that the *temporal pressure* represented by time budgets in different formats of the game should further amplify these differences. Indeed, an idea very dear to the folklore of chess is that good players do not calculate more, just calculate better. David Bronstein, arguably one of the most inventive chess players, was an adept of this view: “*I have always defended playing under time pressure, and I do not think a shortage of time is a bad thing. On the contrary, I have always thought that fast playing is a measure of the ability to play chess*” (Bronstein and Fürstenberg, 1995).

The reduction of chess expertise to speed is, however, overly simplistic: firstly, there is substantial evidence that chess experts do not search “wider”, they do search “deeper” than weaker players (Holding and Reynolds, 1982; Saariluoma, 1990); secondly, as players are forced to play faster, their ability during regular play under normal time controls becomes less predictive of their performance (Van Der Maas and Wagenmakers, 2005). Not surprisingly, even grandmasters make more errors and blunders under conditions in which they have less time than usual to select their moves (Chabris and Hearst, 2003). Thus, time pressure provokes a selective enhancement of rapid object recognition, favoring the

best players, but also increases the likelihood of errors and blunders, which in turn tends to equalize the game. Time constraints and playing ability therefore interact in a highly non-trivial manner.

The consideration of time, therefore, introduces a wealth of dimensions to evaluate the complexity of the game. Among these possibilities, we concentrate on three broad questions pertaining to the statistical structure of chess play under time constraint: What *dynamic* features characterize the evolution of the game, and the strategies of stronger (as opposed to weaker) players in particular? What actual form does the expected trade-off between time and execution accuracy has? Can we assert that the game is a “closed system”, whose dynamics are emergent features of the interaction between the players (as opposed to being solely determined by the *a priori* difference in expertise)?

Given the intricacies of the game, a robust statistical answer to these queries requires a solid experimental framework designed to provide large datasets. Among the various game formats, *rapid chess* provides an unparalleled laboratory to understand decision making in a natural environment. In every game of rapid chess players make around 40 movements, each comprising a decision. The total time budget is finite (a total of 3 min for the games studied here) and hence players need to adopt (implicitly in most cases) a specific policy of time usage. A fundamental advantage of this setup is that a measure of the outcome of each decision can be determined accurately. Moreover, the level of play of both players is well defined. The most relevant aspect of this cognitive experiment is, however, the amount of data it produces: using web-based conduits¹, thousands of players play simultaneously, making millions of decisions per day that can be easily recorded.

¹<http://www.freechess.org/>

Here we have used rapid chess as a laboratory to explore decision making in a natural setup. We have studied the structure of the time players take to make a move during a game, and analyzed millions of instances. This approach allowed us to identify a number of statistical fingerprints that uniquely characterize the emergent structure of the game.

Our results revealed consistent landmarks of the statistics of response times (RTs) distributions and of the fluctuations in score: RTs are long-tailed and vary widely throughout the game. RTs are also highly correlated, when considering the times of a unique player throughout the game and also when considering the correlation across opponent players, indicating that play is not dictated by a state-function. These have theoretical implications, since they question assumptions in models of sequential decision making (Littman, 1996), hence questioning their validity to describe human chained decision making. Our findings also have practical implications: (1) We show that while blunders are more typical in weak players, this dependency is modest and insufficient to classify a player (in rapid chess) based on the number of blunders. (2) The capability of determining an empiric probability function of winning likelihood combining time and score. (3) Our findings may serve to inspire computer algorithms – and more generically game playing algorithms where a sequence of an unpredictable number of moves has to be made with a finite time budget – based on a measure of the efficiency of different time policies.

MATERIALS AND METHODS

DATA ACQUISITION

All games were downloaded from FICS (Free Internet Chess Server)², a free ICS-compatible server for playing chess games through Internet. This server is on-line since 1995, and has more than 300,000 registered users. Each registered user has associated a rating that indicates the chess skills strength of the player, represented by a number typically between 1000 and 3000 points. The rating is a dynamic variable which is updated after each game played according to the Glicko method³. Also, a rating deviation (RD) is used to determine the stability of the rating measure and hence, how much a player's current rating should be trusted. A high RD indicates that the player may not be competing frequently or that the player has not played many games yet at the current rating level. A low RD indicates that the player's rating is fairly well established. For this work, we did not take into account the RD.

Registered players may be humans or computers. These two types of players are distinguished in the server. Only 5% of our database players are computers: 2067 out of a total of 44069 players. For this work (except the analysis reported in the supplemental figure) we discarded all games played by at least one computer.

Users connect to FICS using graphical interfaces, e.g., BabasChess⁴ or Xboard⁵, or command line clients, e.g., telnet. Once connected, users can create, play, and observe games. We developed an application that connects to the FICS server every 30 min and

downloads stored games. The server stores only the last 10 games of each player. We chose to download only the games played by players in the 15% ratings percentile of logged-in players. The application consists of a Python script that connects to the server using TCP sockets and downloads the stored games into a PostgreSQL⁶ database using the standard ICS-compatible instructions (history and smoves). First, the history of logged-in players is queried, checking for repeated games. Once a new game has been found, the detail of game is queried and server answer is parsed and converted into the PGN file format; finally, the PGN is stored in the database. We store the nickname of the logged-in players, the game information (total time, increment, white and black nicknames, players' rating, date, opening variant and result) and the moves of the game with the corresponding time between moves (in milliseconds precision).

Our robot started functioning in May, 2009, downloading only lightning and blitz games, which means total times going from 1 to 15 min. On January 2010 the database consists of more than 2.8 M games (downloading between 10 K and 20 K games per day), resulting in more than 200 M total moves. This is equivalent to a person who played 3 min games for 27 years without leaving the computer.

The participants registered to play in the website are identified by their login name, not their full name, and agree to have their matches stored in a publicly accessible server. The website is designed so that any person, and not just registered participants, can look up the matches browsing the it as a guest. That is, the data are already anonymized. Moreover, in our acquisition process further anonymized the data by stripping all information except the player's ranking. In consequence, individual consent was not sought because of this double layer of anonymity, along with the public, open nature of the website.

A.2 SCORE

An ideal evaluation function would assign to each position three possible values according to the result following best play from both sides: 1, if white is won, 0 is the result is a draw and -1 if black is won. An ideal evaluation function exists for other type of games, as checkers, which is known to result in a draw with perfect play (Schaeffer et al., 2007). However, such ideal evaluation function does not exist for chess and most likely will never be computed according to many theoretical thinkers such as Shannon (1950).

An evaluation function in chess approximates an ideal one considering material value along with other factors affecting the strength of each side. When counting up the material for each side, typical values for pieces are 1 point for a pawn, 3 points for a knight or bishop, 5 points for a rook, and 9 points for a queen. The king is sometimes given an arbitrary high value such as 200 points (Shannon, 1950), or any other value which adds more than all the remaining factors. Evaluation functions also consider factors such as pawn structure, the fact that a pair of bishops are usually worth more, centralized pieces are worth more, and so on. All these factors are collapsed on a single scalar, the score, typically measured in hundredths of a pawn, which provides an integral measure of the goodness of a position. Then, the evaluation is a continuous function which assigns a score (often also referred as value) to

²<http://www.freechess.org/>

³<http://www.glicko.net/glicko.html>

⁴<http://www.babaschess.net/>

⁵<http://www.gnu.org/software/xboard/>

⁶<http://www.postgresql.org/>

each position that estimates of the likelihood of the final result. Conventionally, positive values indicate that the most probable outcome is a win for white.

If after a white move, the score drops abruptly, white winning chances decrease correspondingly. It is then said, in the chess jargon, that white has blundered. Seemingly, if the result of a black move is that the score goes up abruptly, he has lost winning chances, committing a blunder. Hence, the measure $\Delta S = [S(i+1) - S(i)] \cdot C$ where C is the color function (-1 for black moves and $+1$ for white moves) provides a measure of the goodness of the move. Negative (positive) values of ΔS indicate that the moving side has lost (increase) its winning chances.

We quantify the goodness of a move by calculating each move score. The score is a number between -999 and $+999$, that show the valuation of a certain move; 0 meaning this moves is not good for white nor black, positive and negative values give advantage to white or black, respectively. For the analysis of the results, we saturated the score larger than 10 and smaller than -10 , as these extreme values are not interesting in the current analysis.

We used crafty⁷ an open source chess engine written by Robert Hyatt⁸, to analyze the moves and calculate the score. The analysis consists of evaluating the decision tree from a given board position, up to a predefined depth of move number; we used analysis with eight moves of depth. A chess game consisting of approximately 100 moves would take at least 15 s to be analyzed on a Intel XEON 2.2 GHz, 2 GB RAM. Due to the amount of downloaded games and moves, score is calculated in parallel on a cluster. Every 3 h, new games downloaded are sent automatically to the cluster⁹ to be analyzed in parallel using 24 nodes. In the following table we show the number of games analyzed so far for each total time (without increment):

Total time (minute)	No.
2	11322
4	21987
10	29591
5	37355
3	91340
1	142141

For this work, we took into account only 3-min games without increment, with more than 10 moves but less than 100 moves. These constraints were applied to filter abandoned or long-ending games.

RESULTS

RESPONSE TIME DISTRIBUTIONS

We first examined the distribution of move durations. In what follows we will use the term RT, typically used in psychological research, for the time taken by a player to make a move since the opponent's prior move. RT distributions for different move numbers revealed a clear and expected trend: RTs were rapid during the first moves (the opening of the game) and the last moves

(the endgame) and significantly longer during the middle game (Figure 1A). Accompanying the variations in the mean, we observed different distributions during the distinct stages of the game: rapid moves (during the endgame and opening) show power-law tail behavior (revealed by a linear dependence in the log-log representation), while the distribution of RTs for middle game moves showed pronouncedly longer tails (Figure 1B).

We investigated whether higher rated players invest their time budget in a different manner, reflecting distinct time policies. We calculated the mean RT of low rated (<1400) and high rated (>1900) players (see Methods for the determination of player rating). High rated players amplify the variations of RTs during the game: they play faster than lower rated players during the opening games, and slower during the middle game (Figure 1C). Interestingly, these effects do not compensate (although this of course depends on the precise definition of the border between middle game and ending, which here is merely approximated) as clearly observed in the cumulative distribution of time used (Figure 1D). For instance, at move 40, higher rated players have used a significantly larger proportion of their time budget than lower rated players (high rated $77.9 \pm 0.3\%$ and low rated $74.7 \pm 0.4\%$). Since the distributions of time budget left is not Gaussian, we performed a Wilcoxon signed rank test comparing the both distributions, which was also highly significant ($z = 5.44$, $p < 10^{-8}$).

In a wide range of experiments, RTs have been shown to covary with the total number of choices, a phenomenon referred as Hick's law (Hick, 1952). Since the pioneering work of de Groot (1965) it was understood that chess was an exception to Hick's law, since players only explore a subset of all possible moves.

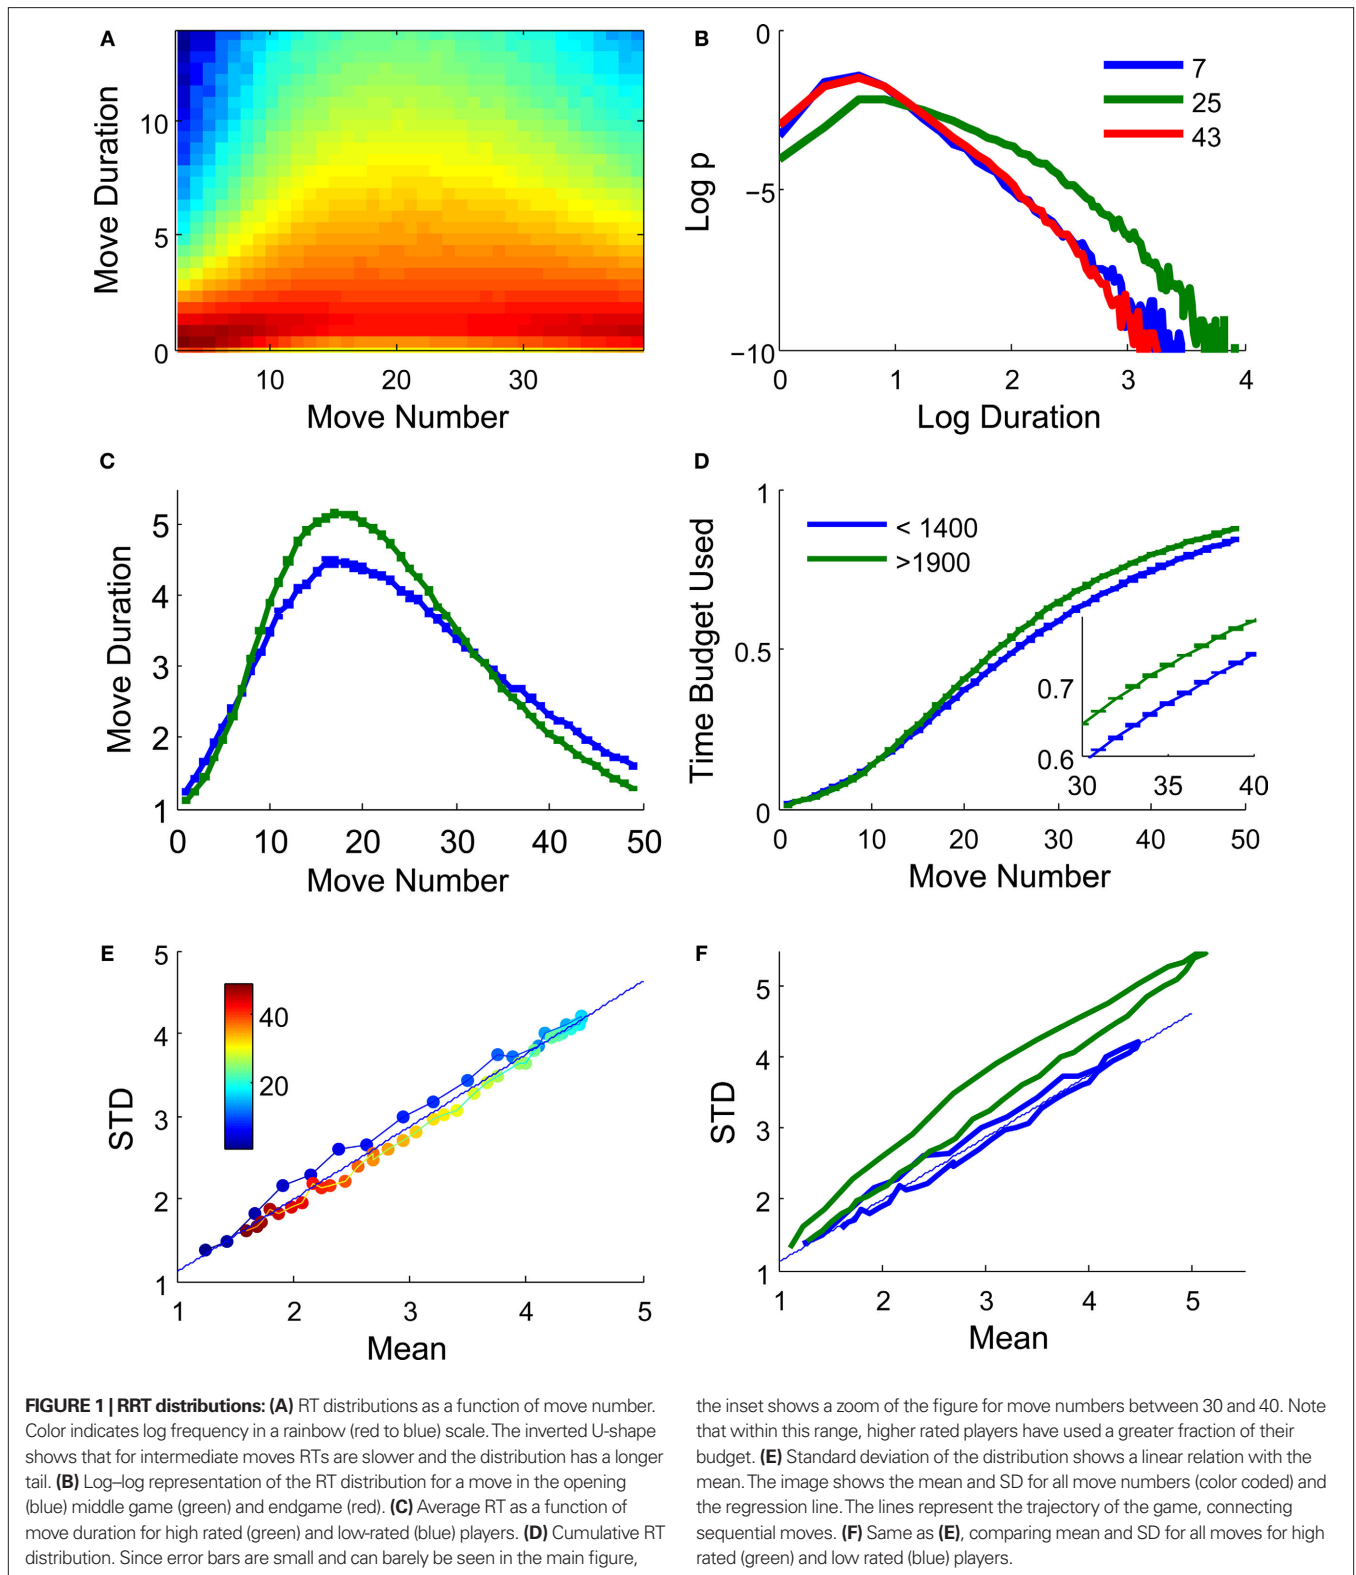
To investigate whether we could partly explain RT variability from an heuristic notion of complexity, we measured a concrete empirical distribution of different moves as a function of move number (Figure 2). For each move number we count – instead of simply all legal moves – the number of different move-types, and their corresponding probabilities, that are actually played in concrete games. A move-type is defined by the piece-moved and the initial and final square. For instance, Nb1–Nc3 indicates that the knight was moved from the square b1 to the square c3.

For the first white move, 20 different moves are possible; advancing each pawn one or two squares and moving each knight to two possible squares. However, not all moves are made with equal frequency, the majority of moves correspond to e2–e4 (advancing the king pawn) or d2–d4 (advancing the queen pawn). As the game progresses, the distribution of move-types becomes broader (Figure 2). This distribution was estimated out of a total of 650,000 games, which assured convergence of these distributions.

We then calculated the entropy of the distribution, which estimates the weighted number of options. The entropy increased with move number, reaching a plateau in the middle game and then decreasing moderately (Figure 2). The decrease in entropy during the endgame is interesting. During endgame the board is less cluttered and hence piece mobility is greater. However, there are less pieces, restricting the number of possible moves. This is clearly illustrated, when observing the probability of piece-move throughout the game (Figure 2).

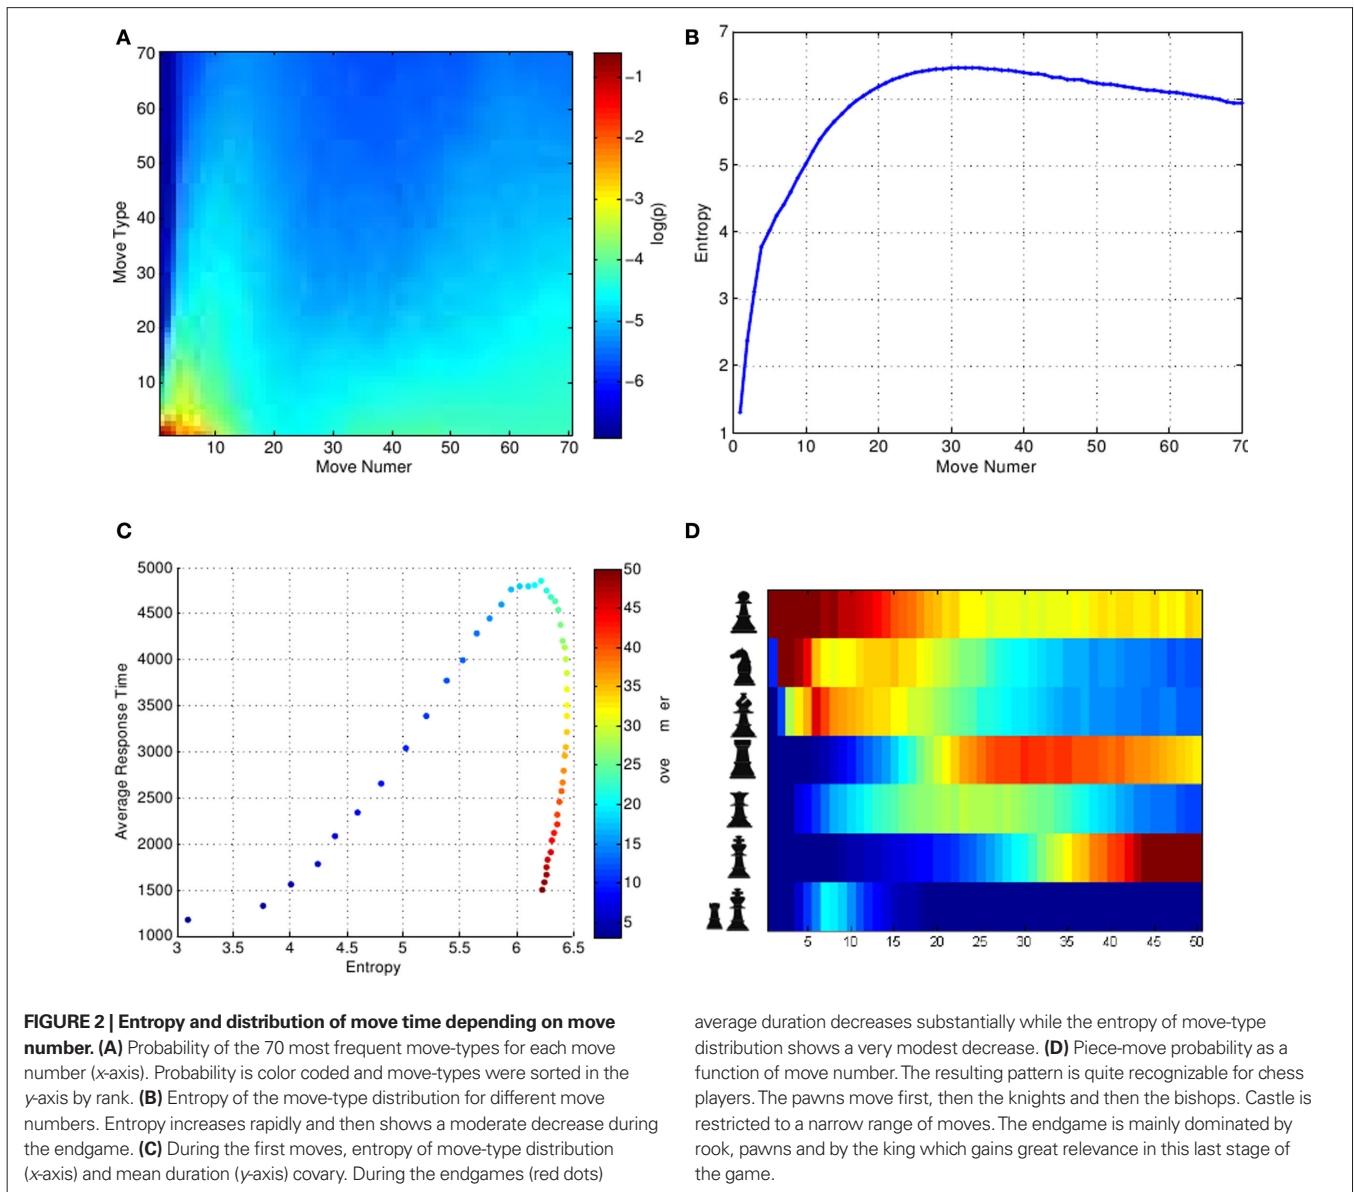
To investigate whether move-type entropy could partially explain the duration variability, we plotted the average move duration as a function of the entropy of move-types distribution (Figure 2).

⁷<http://www.craftychess.com/>
⁸<http://www.cis.uab.edu/hyatt/>
⁹<http://cecar.fcen.uba.ar>



For the first moves, these observables showed a strong correlation, indicating that such broad measure explains some aspects of the trends observed in average duration (Figure 1A). However, for the last moves, the decrease in entropy was very moderate compared

with a major decrease in average move duration. This suggests that speeding of last moves is not dominated by the complexity of the position, being most likely determined by time pressure (Figure 1D).



Finally, we explored whether there was a linear relation between the mean and the standard deviation of the RT distributions, another hallmark of human decision making (Wagenmakers and Brown, 2007). We measured the mean and standard deviation of RT distributions for different move numbers. As observed in laboratory experiments, mean and standard deviation are related linearly (Figures 1E,F). Higher rated players show a greater scaling of variability with the mean (low rated players, regression: $SD = 0.1 s + 0.91 \langle RT \rangle$, for high rated players $SD = 0.6 s + 1.36 \langle RT \rangle$). This further emphasizes that lower rated players play with a more conserved time policy while higher rated players deploy time in a more variable manner, probably depending on necessities of the position.

CORRELATIONS IN RESPONSE TIME DISTRIBUTIONS

Chess should be played as a state-function. In perfect play, the best move is solely a function of the position and should not depend on how this position was reached or in other elements of

the history of the previous moves¹. In practical terms, however, it is not expected that a human player will play according to a state-function. Rather, the existence of plans and schemes, the influence of previous thoughts and evaluations, the assessment of the opponent ability in deciding whether to gamble on a risky move, and in particular, the presence of a time constrain, are likely to be reflected in an interaction between successive moves.

Here we explored quantitatively the existence of such interaction, measuring the linear correlation in RT across different moves. As an example to ground this measure, a positive correlation between

¹By position here we mean a set of statements which include a statement of the position of each piece and other statements which provide information of previous moves: whether the king has previously moved or not (to know if castle is possible), the number of moves made since the last pawn move or capture (which has to be known since after 50 such moves the game is drawn), etc.

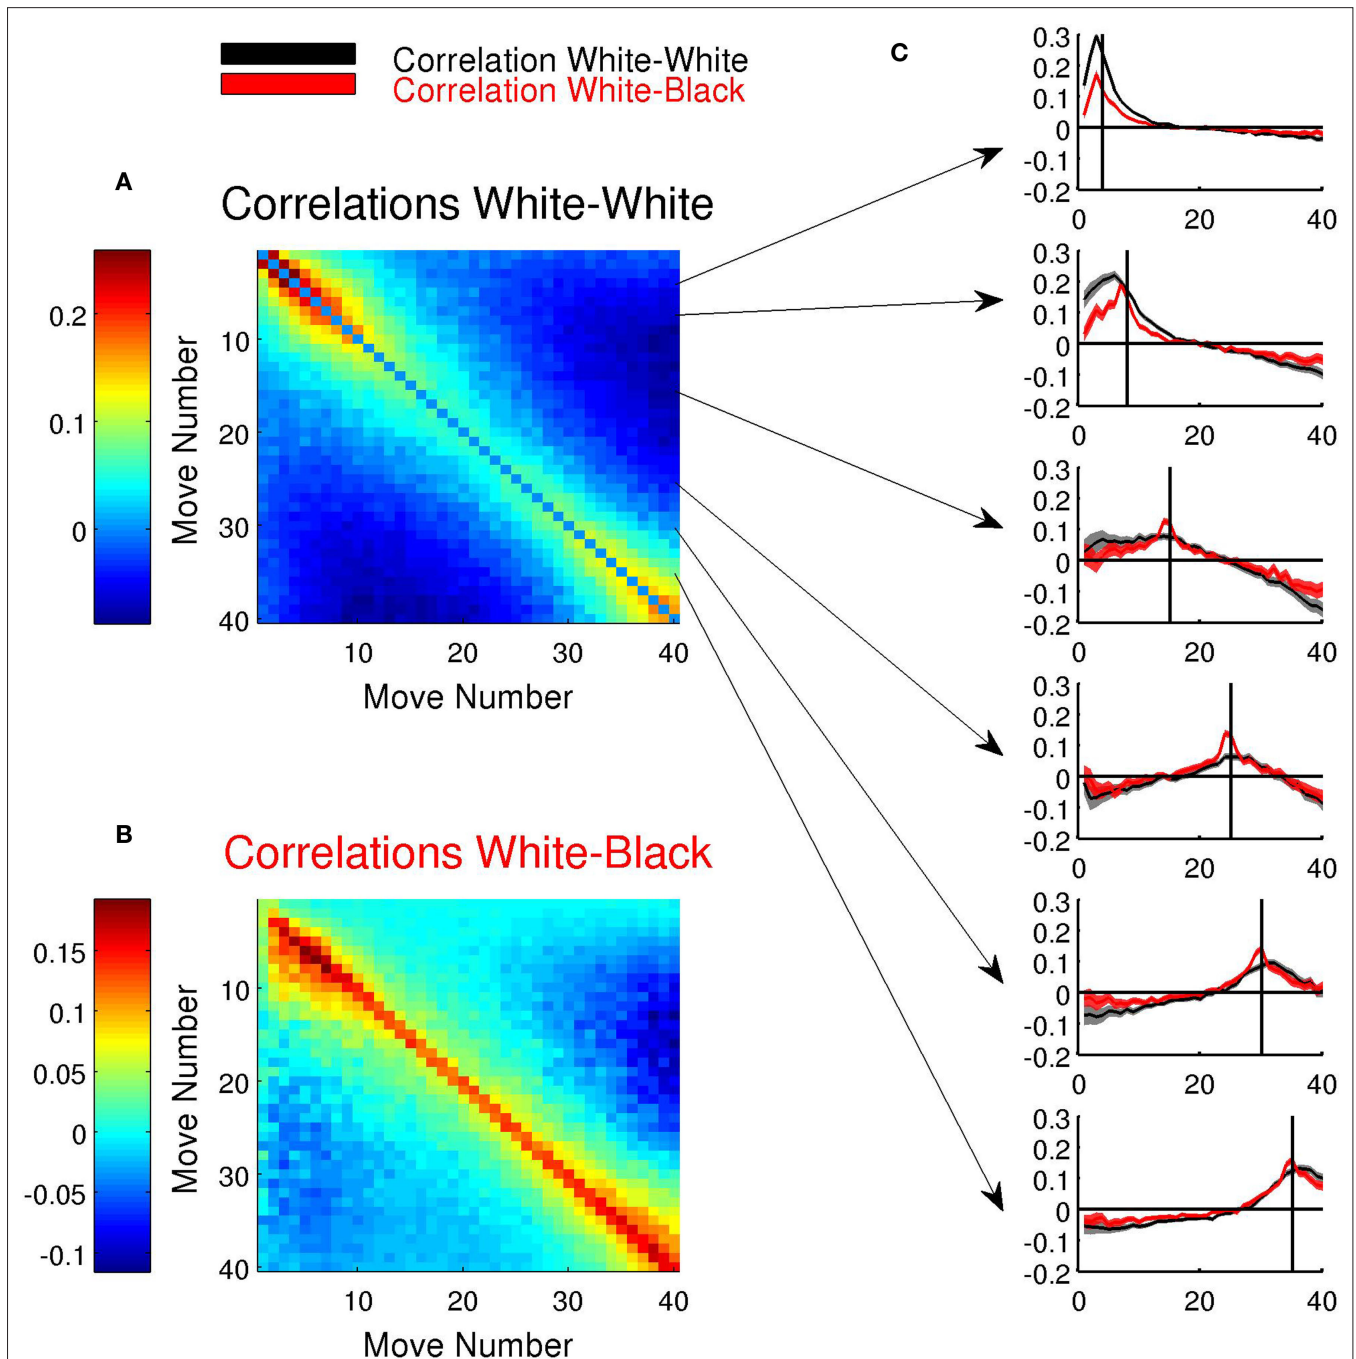


FIGURE 3 | Correlations in RT distributions: (A) Correlation matrix of move durations. The color of the position (i, j) of the image codes the correlation of white-move i , with white-move j . **(B)** Same as **(A)**, comparing the correlation of the duration of players playing white and black.

(C) Detail of the correlation of a specific move with all other moves. The shaded bars indicate the 0.01 confidence interval of the correlation. The vertical line indicates the particular move to which the correlation is plotted.

moves 10 and 11 implies that when a player makes move 10 slower (faster) than the average time of the tenth move, then move 11 will also be played slower (faster) than average.

The cross-correlations amongst all moves revealed a conserved pattern: correlations were positive for consecutive moves and significantly negative for long-range differences in move number (**Figure 3A**).

Significance of the correlations were measured calculating the statistics of a linear regression. For each pair of distributions of moves i and j we performed a linear regression $d^{(i)} = \beta 1_{ij} d^{(j)} + \beta 0_{ij}$ and errors (shadow bars in **Figure 3C**) were estimated calculating the 95% confidence intervals for the coefficient estimates.

Negative correlations were simply expected by the finite time budget. The positive neighboring correlations indicate that the overall dominating trend amongst close moves is to deviate coherently from the mean. As with our other analysis, this broad statistical marker collapses several patterns. Hence, if a position is difficult (and thus requires thinking time) the position of the following move is also likely to be difficult and thus long sequences of slow moves are expected, resulting in positive correlations. A different trend can also be expected due to the existence of plans, or conception of long sequences of chunked moves. Then, a slow move (for instance, when global plan or strategy is being set) is likely to be followed by rapid moves (the practical execution of the plan). Our results merely show that these interaction effects create a positive correlation and provide a measure of the kernel (the number of moves) in which a significant interaction is observed. Interestingly, this kernel varies with move number: correlations in the opening are much stronger and narrow than in the middle game (compare the first and middle panels of **Figure 3B**). This result is consistent with the existence of qualitative different forms of play: rapid opening and slow middle game. A slow opening move may reflect that the player is “out of the book”, i.e., without prior knowledge of the position. This certainly implies that subsequent moves will similarly be performed without opening knowledge and thus comparatively slower. Hence, correlations are markedly positive.

To quantify this observation we measured, for each move number, the kernel size at different thresholds of correlation. When the correlation was set at a relative high value ($C = 0.1$) only moves in the opening and endgame were correlated with subsequent moves. The size of this kernel varied between 5 to 7 for the first 10 moves and for moves after the 35th move. On the contrary moves in the middle game did not correlate with neighboring moves at such threshold. The contrary result was observed when measuring the number of moves with significant correlation at a low threshold ($C = 0.01$). The kernel-size measure showed a peak at the middle game (move 18) and was minimal during the opening and endgame.

The previous correlation analysis explored the interactions of successive moves of a single player. In chess, as most two-player games, the behavior of both players does not need to be independent. Just to guide intuition, we mention two concrete examples from the folklore of chess: firstly, complicated positions are typically complicated for both sides, establishing an inter-player correlation (similarly, simple positions tend to be simple for both sides). Certainly, there are exceptions to this rule; many positions are easy to play for one side and difficult to play for the other. Secondly, a player may use the opponent’s time to think his future moves. To avoid this, players often play fast when the rival has little time left, resulting in positive correlation of move durations of both players. We examined this hypothesis numerically, measuring $\langle d^w(i), d^b(j) \rangle$, the correlation between the duration of the i -th white-move and the duration of the j -th black-move, for all values of i and j (**Figure 3B**).

Black–white correlations showed virtually the same pattern as white–white correlations, indicating that the usage of both clocks is remarkably tied. Note that the max of black–white correlation is greater than that of white–white correlations. This seemingly paradoxical result can be easily understood: $\langle d^b(i), d^w(i + 1) \rangle$ reflects

temporal correlations between moves separated only by one ply (successive movements in the game), while $\langle d^w(i), d^w(i + 1) \rangle$ correlations between moves separated by two plies. A ply (plural plies) refers to one turn taken by one player and, correspondingly, to one level of the game tree. Thus, our results show that the correlation proved to be stronger for single ply separations even when this involved opponent players than for successive moves of the same player.

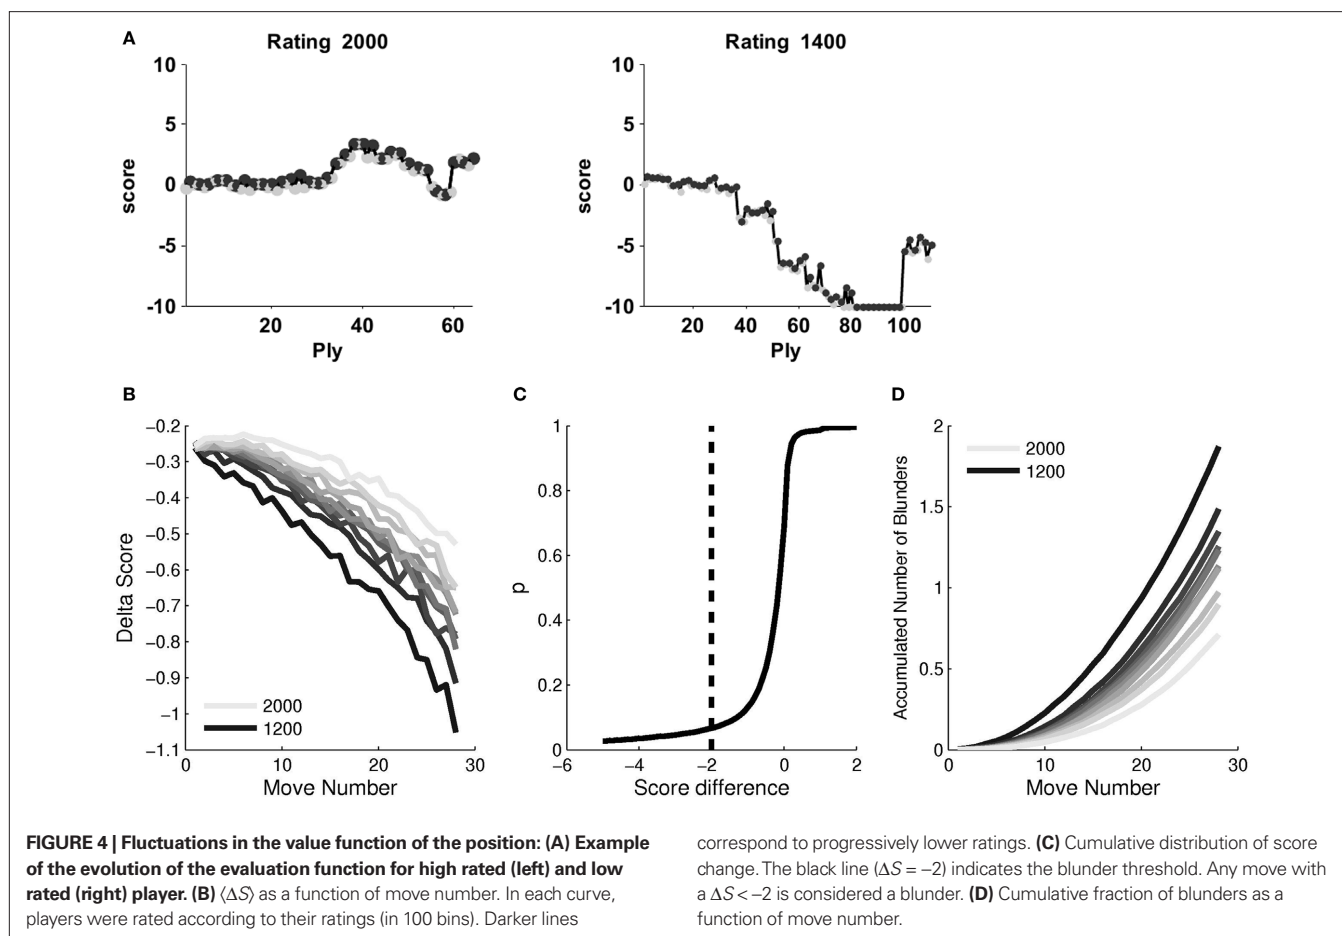
THE DYNAMICS OF SCORE

Human decisions can be classified, as done in the previous section, according to the time they take. They can also be classified according to their outcome, i.e., an objective measure of the goodness of choice. In controlled laboratory experiments, conditions can be set so that the goodness of a decision is under control (by parametrically varying the reward of each choice). In more realistic setups, the outcome of choice is more difficult to quantify, and indirect measurements such as self-reported post-choice satisfaction have been used to estimate the goodness of choice (Dijksterhuis et al., 2006). A great advantage of analyzing decisions in a chess game is that current chess software can rapidly and accurately evaluate a position, providing what it is usually called a *score* for each move number, $S(i)$ (see Methods for details).

Conventionally, positive values of S indicate that the most probable outcome is a win for white. If after a white move, the score drops abruptly, white winning chances decrease correspondingly. It is then said, in the chess jargon, that white has blundered. Seemingly, if the result of a black move is that the score goes up abruptly, he has lost winning chances, committing a blunder. Hence, the measure $\Delta S = [S(i + 1) - S(i)] \cdot C$ where C is the color function (-1 for black moves and $+1$ for white moves) provides a measure of the goodness of the move. Negative (positive) values of ΔS indicate that the moving side has lost (increase) its winning chances.

The score of a typical game of chess is a fluctuating variable. In some games one side takes the advantage and remains with a better position until the end in which case the score function does not change sign. In some games, the advantage progresses slowly and in other games it can change abruptly. The evolution of score of representative games of high rated and low rated players can be seen in **Figure 4A**). As intuition suggests, the fluctuations in a game between low rated players are larger than in a game between high rated players. Intuitively, one would also expect that fluctuations in score should increase with move number. During opening play, players are less likely to blunder less since they use previous knowledge and less pieces are exposed. In addition, as the game proceeds, players have progressively less time left and thus the resulting moves are more prone to errors.

To examine these two predictions we measured $\langle \Delta S \rangle$ as a function of move number and player rating (**Figure 4B**). As expected, we verified that, for players of all ratings, $\langle \Delta S \rangle$ decreased monotonically with move number. In addition, rating had also a significant effect in $\langle \Delta S \rangle$ which was larger for lower rated players than for higher rated players for all move numbers. These two effects showed an interaction: while for the highest rated players $\langle \Delta S \rangle$ was constant during the first 10 moves and decreased thereafter, for the lowest rated players $\langle \Delta S \rangle$ decreased steadily from the first moves of the game (see Methods for a technical explanation of this trend).



We generated a grand distribution of $\langle \Delta S \rangle$, collapsed across all moves in the database and measured the cumulative histogram (Figure 4C) which revealed that the $\langle \Delta S \rangle$ is clearly asymmetric and biased toward the negative values.

Understanding the significance of this observation requires some understanding of the way the engine's evaluation function works. At any given position, the score is determined following a finite tree driven exploration of successive moves. In our case the depth of the search was set to eight moves (16 plies). This signifies that, when the engine assigns a score of s to a position, it considers that this is the score of the resulting position when the best moves have played by both sides. According to the engine's exploration, other moves would worsen the score. Hence, when a player makes a move which results in a positive value of $\langle \Delta S \rangle$ – i.e., the value of its position has increased according to the engine – it evidences the non-perfect play of the engine. This may origin from two different reasons: first, simply, when the move has been made, the same depth advances further in the game which may improve the assessment of the position. Alternatively, the player may have chosen a move which the engine had not even considered. In either case, it shows that the engine is not providing a perfect evaluation of the position. Hence, the fraction of moves with a positive value of $\langle \Delta S \rangle$ provides an estimate of the goodness of the engine relative to the players of the database.

Indeed, only 0.5% of the moves had a value of $\langle \Delta S \rangle$ greater than a standard deviation of this distribution (1.56). On the contrary 8.3% of the moves (a 16-fold increase) had a value $\langle \Delta S \rangle < -1.56$. This asymmetry is even more pronounced if one does not consider moves from the endgame (where the engine is likely to underestimate the score of a position and hence to improve its evaluation as the game moves on; for instance, as a pawn gets closer to promotion): 0.2% of the moves had $\langle \Delta S \rangle > 1.56$ and 10% (i.e., a 50-fold increase) $\langle \Delta S \rangle < -1.56$. This result testifies the quality of the evaluation engine relative to the examined players in our database.

Finally, we converted the analogous $\langle \Delta S \rangle$ into a discrete measure of blunders, defining a blunder as a move whose $\langle \Delta S \rangle < -2$. From the cumulative $\langle \Delta S \rangle$ histogram it is quite clear that the results are not strongly dependent on small perturbations of this threshold. We then measured the number of accumulated blunders as a function of move number and rating. Low rated players make (on average) their first blunder before move 17. High rated players do not make (on average) a blunder until move 40 (Figure 4D).

SCORE AND TIME: THE PRICE OF A SECOND

In most psychological experiments, a relation is found between the duration of a choice and its outcome. The most widely studied example is the speed-accuracy trade-off (Corbett and Wickelgren, 1978) according to which participants may invest more time to assure more accurate performance. These factors may interact

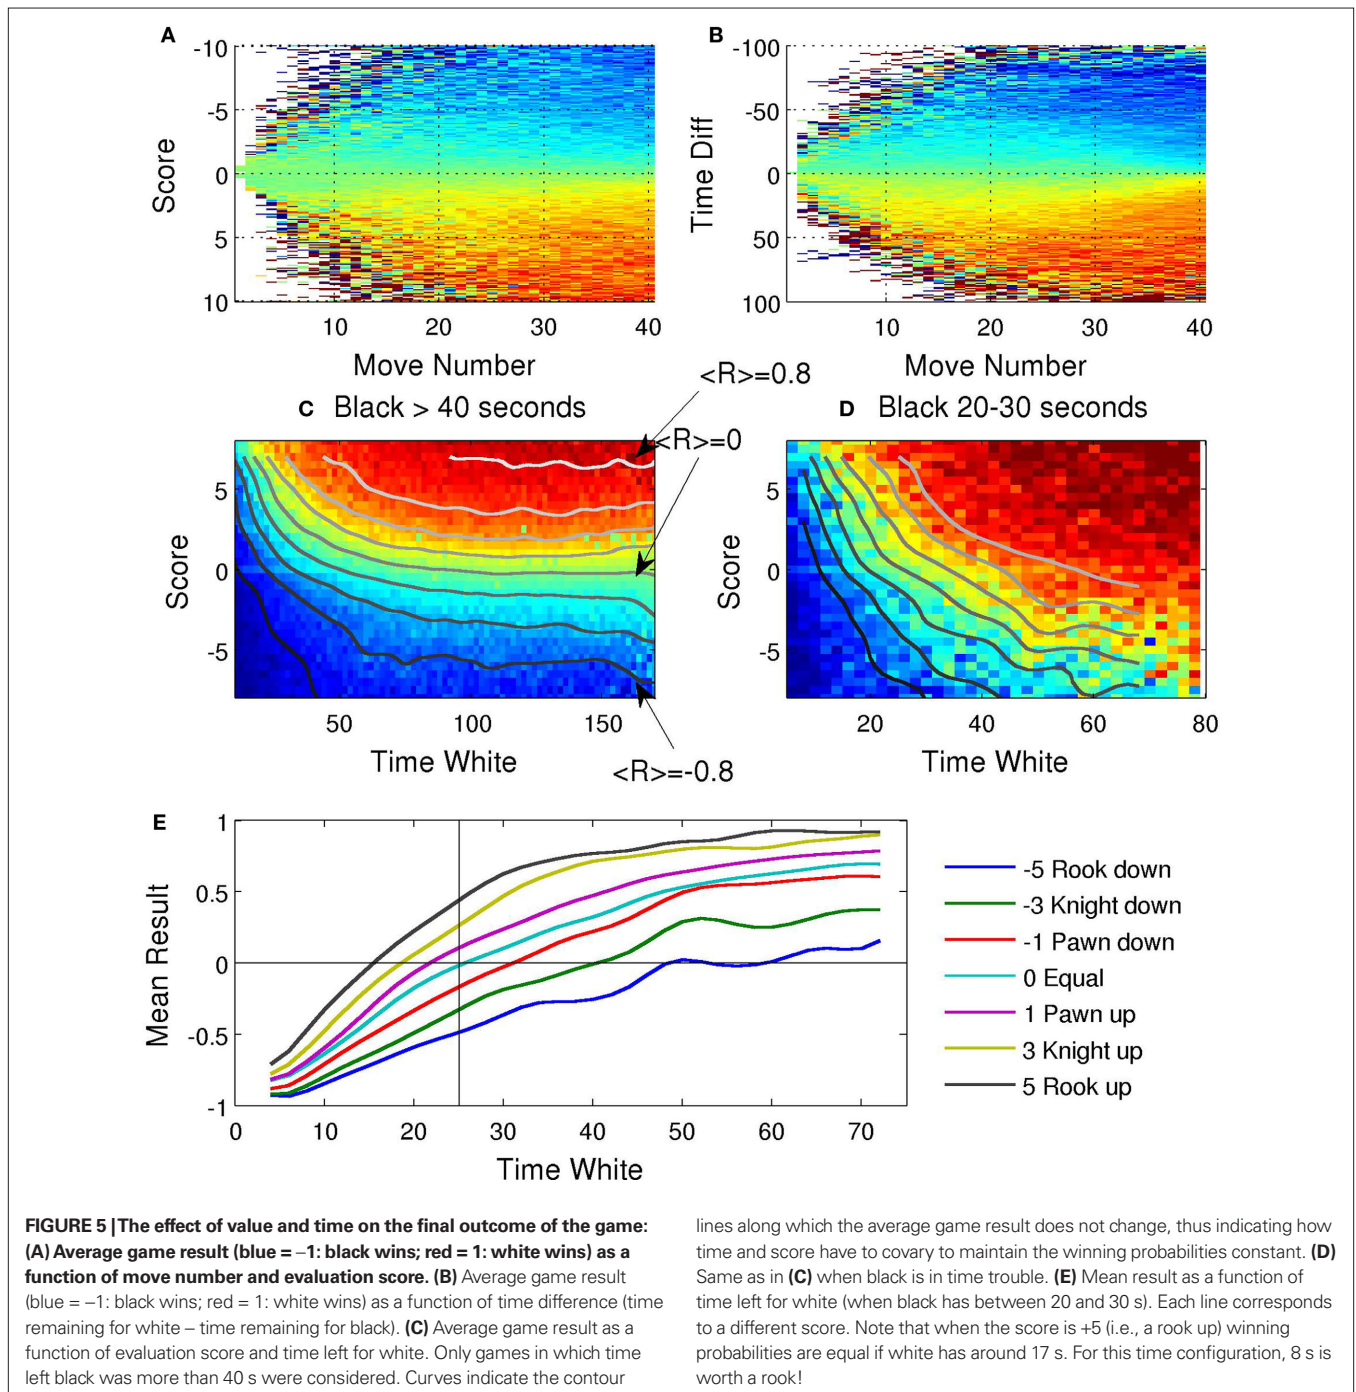
following more complicated rules when the “complexity” of the decision, the urge to respond and many other factors of the decision may vary in a correlated manner (Gold and Shadlen, 2002; Kamienkowski and Sigman, 2008; Zylberberg et al., 2010).

In games with finite time, players adopt (typically implicitly) a strategy compromising accuracy and speed. Some chess players are known to use the clock “badly”, leading often to situations in which the over-evaluation of subtle differences in a position leaves them without time for a significant portion of the game. The findings presented in the first section revealed clear markers of regular

time policies (fast play in the opening moves and slow play in the middle game), suggesting that there is a correlation between chess rating and policy choice.

In this section we sought to measure the combined effects of time and score, to estimate the price of a second measured in pawns, knights, and open files². As described previously, the score is a con-

²Time here refers to the clock time, measured in seconds. Just to avoid misleading we emphasize that here we do not refer to the notion of initiative or development which is also referred as a time factor often in the chess game.



tinuous estimate of the winning probability of each side. **Figure 5A** simply testifies this known fact, displaying the average games in which white won as a function of the score, for different move numbers.

For this analysis we grouped moves independently of the time left. **Figure 5B** illustrates the complementary analysis. It shows the average number of games in which white won as a function of the difference in time left. As expected, the figure shows that the side which has more time at its disposal has a greater probability of winning. This finding begs the question as to how these factors combine. In other words, given the triplet $\{T_{\text{white}}, T_{\text{black}}, S\}$, is it possible to determine the side with greater winning probability?

We next concentrate on some critical behavior of this probability function, illustrated in **Figure 5C**. We represent the winning probability (color coded, red indicates probability 1 of white winning and blue probability 1 of black winning) as a function of T_{white} and S . This analysis was restricted to moves for which black had sufficient time, i.e., more than 40 s, at least about 1/4 of their original time budget. The limit behaviors are quite evident: when white has virtually no time (left portion of the image) his winning chances are close to 0, even when it has a better position. When both sides have sufficient time (right portion of the image), winning probability is almost exclusively a function of the score. We plotted the contour lines of the function, i.e., the curves in the $T_{\text{white}} - S$ plane along which the winning probability for white is constant. Contour lines are almost flat for large time budgets. Changing T_{white} from 150 to 50 s is compensated by a small fraction of a pawn in the score function. In contrast, when the total time of white becomes shorter, the contour lines become very steep. In this regime (see also **Figure 5D**), a mere few seconds can change significantly the winning probability function, largely overriding score differences.

To quantify this observation we estimated the price of time on time trouble situations. We considered all moves in which black had between 20 to 30 s remaining. For each values of S , we measured which value of T_{white} resulted in equal winning probabilities for both sides. For $S = 0$, the winning probabilities equalize when white has on average the same time than black (25 s); this simply serves as a control. We then determined the value of white time left, at which probabilities equalize when the score difference is +3 (roughly when white is a minor piece up without other forms of compensation). Analysis revealed that this value is 17 s. Thus, a small difference of 8 s (note the difference in time between the crossing of the cyan and green lines, **Figure 5E**) is on average sufficient to compensate for a full piece when the time left is between 20 to 30 s. This function is clearly asymmetric. If black is on the advantage of a clear minor piece, then winning probabilities equalize when white has (to compensate) a total of 40 s at his disposal, 15 more seconds than black. When these data are grouped on fractional quantities, the time difference measured relative to the total time, the two estimates become comparable indicating that, within this degree of time trouble (about half a minute of total time to complete the game), a 50% increase in total time is roughly equivalent to a minor piece.

DISCUSSION

Here we have studied the statistics of a large ensemble of rapid chess games, concentrating on questions of time usage, value as a measure of execution efficacy, and the relationship between them as a function of the game's progression. Our results revealed highly reliable patterns:

1. RT distributions are heavy-tailed and show a qualitatively distinct shape at different stages of the game.
2. RTs are positively correlated for successive moves of both players. For more distant moves, as expected for a finite time budget, correlations are negative.
3. Fluctuations of ΔS increase steadily throughout the game. Fluctuations decrease for higher rated players but not in a very sharp manner.
4. The winning likelihood $\langle R \rangle$ can be reliably estimated from a combination of time and score. With sufficient time (more than a minute), changes in score are not compensated by time saving. When the remaining time is less than 30 s, in practical terms, rapid moves are more efficient than slow moves even if they result in an immediate significant loss of value.

While endgame and opening stages result in comparable RT distributions, these may originate in different causes. Opening play is largely determined by memorized schemes leading to very rapid responses. Endgame RTs are also very rapid which may result from the fact that, at this stage of the game, the time budget has largely been spent and thus players have to respond very fast. In certain cases (for instance in a Rook + King vs King ending) players also play at great speed using a memorized sequence or algorithm. These different sources of RT variability are blurred in our analysis which simply captures the main emergent statistical elements.

A large body of research has studied human RT in laboratory setups, reliably finding that they systematically result in skewed heavy-tailed distributions for different experimental paradigms (Luce, 1986). In particular for chess, it has been previously reported that RT distributions are non-Gaussian (Van Der Maas and Wagenmakers, 2005). We find that, more precisely, the distributions display a power-law behavior over several decades, a feature that has several theoretical implications. It implies that a wide spectrum of time scales are, if not equally, at least significantly involved in the dynamics of such a process, and likely interacting with each other. The lack of defined scales has been postulated as a hallmark of emergent behavior in complex systems by Bak (1996) in the theory of self-organized criticality and places a clear constraint on the modeling of chess as a decision process. Crucially, the finite time budget poses a definite cutoff for power-law distributions of RT in complex decision making. Buridan's metaphor – wherein an ass, placed precisely midway between a stack of hay and pail of water, dies of both hunger and thirst since it cannot make any rational decision to choose one over the other – of infinite RT in even decisions may not be out of relevance. Indeed, single chess decisions can scale up to hours indicating that, in the absence of a budget forced cutoff, power-law behavior may extend over several more decades. An important follow up of this study will be to determine the rescaling of RT distributions when the total time budget changes.

The effect of chess expertise on RTs has also been studied using classic psychological tasks (change blindness, Stroop, detection, ...). Reingold et al. (2001b) studied the effect of distractors on RTs in check detection tasks. They found that distractors had no effect in expert players indicating parallel and automatic processing of the board in expert play. Further supporting this view, in an extensive program investigating eye-movement in chess problems, they showed that

the perceptual advantage demonstrated by chess experts is mediated by a larger visual span for chess-related, but not for chess-unrelated visual patterns (Reingold et al., 2001a; Reingold and Charness, 2005). Our approach, relying on massive sampling, allowed us to investigate the structure of RTs during concrete play.

The correlation structure observed in RT data also constrains the possible models which may account for this data. First, it discards at once any model in which decisions are made independently, or even in which decisions are made only taking into account the remaining time budget. The strong positive correlations simply discard this possibility. It remains also for further studies to unfold the many different contributions to RT interactions which integrate in our measure of correlation: first, a player cannot avoid being primed or influenced by the inertia of previous thoughts (Shallice, 1982, 1988; Shallice and Burgess, 1996). Chess players know that this is a well known cause of blunders: in their calculations they often confuse pieces positions with ones occupied previously or even in mental variation which were never settled on the board. Also, as the game progresses, a player reads the opponent and sets a risk policy according to an estimation of the opponent abilities or strength. If during the course of a game a player has estimated that his rival is not very strong he may be more willing to take risks. This strategy is very frequent in backgammon, where optimal play depends on risk policies that take the opponent's strength into account. However, the most likely source of correlations is determined by a relative continuous function of complexity of the board. We currently lack a good measure of "chess complexity". As de Groot had identified in his early work, simply counting the number of possible moves, or even the number of good moves does not constitute a good estimator of complexity (de Groot, 1965). A rich indexed database, in which each position is labeled according to all conceivable elements of the chess lexicon (open files, threats, mobility of pieces, number of pieces, time left, ratings, etc.) could be used to conduct regression analysis to measure – by brute force – what elements of a position determine that a player will spend a long time sitting on it before making a move.

A class of algorithms, broadly known as Markov decision processes (Littman, 1996) have been very useful and widely studied in artificial intelligence to model sequential decision making. These models generically assume that decisions are made in a stationary environment and are state-functions. The results described here challenge the validity of this class of theoretical models to understand human decision making. First, the observation of heavy-tailed RT distribution in time-constrained human decision making imply that long time scales are statistically significant, essentially violating the stationarity premise. Second, the observation of long-range positive correlations in RT (unrelated to finite time budget) violate the state-function hypothesis. Our finding, therefore, will require the development of new theoretical models to accommodate decision making in non-stationary, transient settings. It also raises an interesting series of questions to understand the discrepancies between formal models with great practical relevance (Littman, 1996) and human-computation inspired algorithms.

The steady increase in score fluctuations is an expected feature given the time limitation rule in rapid chess. This is the case for players of all ratings, even though there are clear differences between the lower and higher rate tiers. It is less obvious, though, to note that there are no sharp transitions or non-linear separations in

the blunder behavior as a function of the player's rate. In a way, this implies that rapid chess, while demanding, still has the ability to tap into cognitive functions that do not admit over-training or high specialization. This interpretation is also well aligned with our finding that when the remaining time is relatively short, it is more effective for a player to be able to make a decision, *any* decision, than to ponder several alternatives. The finding that blunder count only modestly separated high and low rated players also motivates further research in this direction. A possible improvement of the method is to classify blunders not only based on their score difference but also on the "complexity" of the blunder. For instance, a move that loses a piece because it simply puts it in a position where it can be captured without cost, is a very "evident" blunder not expected in high rated players. On the contrary, a move that loses a piece because after a complex combination involving a long sequences of forced-moves it is found that the piece is lost, constitutes a more subtle blunder which may be more typical in high ranked players. At an analysis of high depth both moves are equally ranked as blunders with the same score difference. In practical terms, such an analysis may involve including analysis depth, as well as score, to catalog blunders as an efficient manner to rank players. It will also be useful, in this context, to understand the patterns of score and blunder behavior in computer vs computer and computer vs human games, as a means to further constraint models of human decision making under temporal constraints.

The proliferation of chess servers on the Internet has turned active chess, blitz and lightning, into a vast cognitive phenomenon involving engaged participants. This large database of human decision making can be used as a privileged window to understand human cognition. Computer scientist have been recently embarked in a project to direct voluntary use of human computing cycles in a coherent and productive direction (Von Ahn and Dabbish, 2004; Von Ahn, 2006; Von Ahn et al., 2006, 2008). For instance, in GWAP (Games With a Purpose) people play a game in which they determine the contents of images by providing meaningful labels for them. Thus, a computationally intractable problem (image labeling) is solved by encouraging people to do the work by taking advantage of their desire to be entertained. Here, we use in a similar vein a leisurely cognitive activity, rapid chess, as a window into cognition. We hope that our work may prompt other large-scale studies in chess as well as similar decision making activities. As Gary Kasparov suggests in his book "How Life Imitates Chess: Making the Right Moves, from the Board to the Boardroom" chess is more than a metaphor: it makes the case for using chess as a model for understanding and improving human decision making everywhere else.

ACKNOWLEDGMENTS

We thank Robert Hyatt for developing Crafty as an open source project and for having promptly responded to many questions which have greatly contributed to the development of this work. We also thank the developers and administrators of FICS for continuously supporting a free chess server of the highest quality and completely open to experimentation. This work was supported by the Human Frontiers Science Program. The computing power was partially provided by Centro de Computación de Alto Rendimiento (CeCAR), <http://pme84.exp.dc.uba.ar/>.

REFERENCES

- Bak, P. (1996). *How Nature Works: The Science of Self-Organized Criticality*. New York: Copernicus.
- Bronstein, D., and Fürstenberg, T. (1995). *The Sorcerer's Apprentice*. London: Cadogan Books.
- Burns, B. (2004). The effects of speed on skilled chess performance. *Psychol. Sci.* 15, 442–447.
- Chabris, C., and Hearst, E. (2003). Visualization, pattern recognition, and forward search: effects of playing speed and sight of the position on grandmaster chess errors. *Cogn. Sci.* 27, 637–648.
- Charness, N. (1981). Visual short-term memory and aging in chess players. *J. Gerontol.* 36, 615.
- Corbett, A., and Wickelgren, W. (1978). Semantic memory retrieval: analysis by speed accuracy tradeoff functions. *Q. J. Exp. Psychol.* 30, 1.
- de Groot, A. (1965). *Thought and Choice in Chess*. The Hague: Mouton.
- Dijksterhuis, A., Bos, M., Nordgren, L., and Van Baaren, R. (2006). On making the right choice: the deliberation-without-attention effect. *Science* 311, 1005.
- Gobet, F., and Simon, H. (1996a). Templates in chess memory: a mechanism for recalling several boards. *Cogn. Psychol.* 31, 1–40.
- Gobet, F., and Simon, H. (1996b). Recall of random and distorted chess positions: implications for the theory of expertise. *Mem. Cogn.* 24, 493.
- Gold, J., and Shadlen, M. (2002). Banburismus and the brain decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299–308.
- Hick, W. (1952). On the rate of gain of information. *Q. J. Exp. Psychol.* 4, 11–26.
- Holding, D., and Reynolds, R. (1982). Recall or evaluation of chess positions as determinants of chess skill. *Mem. Cogn.* 10, 237–242.
- Kamienkowski, J., and Sigman, M. (2008). Delays without mistakes: response time and error distributions in dual-task. *PLoS ONE* 3, e3196. doi:10.1371/journal.pone.0003196.
- Littman, M. (1996). *Algorithms for Sequential Decision Making*. Ph.D. thesis, Brown University, Providence, RI.
- Luce, R. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. New York: Oxford University Press.
- Reingold, E., and Charness, N. (2005). “Perception in chess: evidence from eye movements,” in *Cognitive Processes in Eye Guidance*, ed. G. Underwood (Oxford: Oxford University Press), 325–354.
- Reingold, E., Charness, N., Pomplun, M., and Stampe, D. (2001a). Visual span in expert chess players: evidence from eye movements. *Psychol. Sci.* 12, 48–55.
- Reingold, E., Charness, N., Schultetus, R., and Stampe, D. (2001b). Perceptual automaticity in expert chess players: parallel encoding of chess relations. *Psychon. Bull. Rev.* 8, 504.
- Saariluoma, P. (1990). “Apperception and restructuring in chess players problem solving,” in *Lines of Thought: Reflections on the Psychology of Thinking*, Vol. 2. eds K. J. Gilhooly, M. T. G. Keane, R. H. Logie, and G. Erdos (New York: Wiley), 41–57.
- Schaeffer, J., Burch, N., Björnsson, Y., Kishimoto, A., Muller, M., Lake, R., Lu, P., and Sutphen, S. (2007). Checkers is solved. *Science* 317, 1518.
- Schultetus, R., and Charness, N. (1999). Recall or evaluation of chess positions revisited: the relationship between memory and evaluation in chess skill. *Am. J. Psychol.* 112, 555–569.
- Shallice, T. (1982). Specific impairments of planning. *Philos. Trans. R. Soc. Lond., B, Biol. Sci.* 298, 199–209.
- Shallice, T. (1988). *From Neuropsychology to Mental Structure*. New York: Cambridge University Press.
- Shallice, T., and Burgess, P. (1996). The domain of supervisory processes and temporal organization of behaviour. *Philos. Trans. R. Soc. Lond., B, Biol. Sci.* 351, 1405.
- Shannon, C. (1950). Programming a computer for playing chess. *Philos. Mag.* 41, 256–275.
- Van Der Maas, H., and Wagenmakers, E. (2005). A psychometric analysis of chess expertise. *Am. J. Psychol.* 118, 29–60.
- Von Ahn, L. (2006). Games with a purpose. *Computer* 39, 92–94.
- Von Ahn, L., and Dabbish, L. (2004). “Labeling images with a computer game,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 319–326.
- Von Ahn, L., Liu, R., and Blum, M. (2006). “Peekaboom: a game for locating objects in images,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY: ACM), 55–64.
- Von Ahn, L., Maurer, B., McMillen, C., Abraham, D., and Blum, M. (2008). reCAPTCHA: human-based character recognition via web security measures. *Science* 321, 1465.
- Wagenmakers, E., and Brown, S. (2007). On the linear relation between the mean and the standard deviation of a response time distribution. *Psychol. Rev.* 114, 830.
- Zylberberg, A., Fernandez Slezak, D., Roelfsema, P. R., Dehaene, S., and Sigman, M. (2010). The brain's router: a cortical network model of serial processing in the primate brain. *PLoS Comput. Biol.* 6, e1000765. doi:10.1371/journal.pcbi.1000765. <http://dx.doi.org/10.1371/journal.pcbi.1000765>.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 19 March 2010; paper pending published: 20 May 2010; accepted: 23 July 2010; published online: 07 October 2010.

Citation: Sigman M, Etchemendy P, Fernandez Slezak D and Cecchi GA (2010) Response time distributions in rapid chess: A large-scale decision making experiment. *Front. Neurosci.* 4:60. doi: 10.3389/fnins.2010.00060

This article was submitted to *Frontiers in Decision Neuroscience*, a specialty of *Frontiers in Neuroscience*.

Copyright © 2010 Sigman, Etchemendy, Fernandez Slezak and Cecchi. This is an open-access article subject to an exclusive license agreement between the authors and the *Frontiers Research Foundation*, which permits unrestricted use, distribution, and reproduction in any medium, provided the original authors and source are credited.