Journal of Experimental Psychology: General

Do Not Fear Your Opponent: Suboptimal Changes of a Prevention Strategy When Facing Stronger Opponents

Diego Fernandez Slezak and Mariano Sigman

Online First Publication, October 17, 2011. doi: 10.1037/a0025761

Journal of Experimental Psychology: General

0096-3445/11/$12.00 DOI: 10.1037/a0025761

Do Not Fear Your Opponent: Suboptimal Changes of a Prevention Strategy When Facing Stronger Opponents

Diego Fernandez Slezak and Mariano Sigman
University of Buenos Aires

The time spent making a decision and its quality define a widely studied trade-off. Some models suggest that the time spent is set to optimize reward, as verified empirically in simple-decision making experiments. However, in a more complex perspective compromising components of regulation focus, ambitions, fear, risk and social variables, adjustment of the speed–accuracy trade-off may not be optimal. Specifically, regulatory focus theory shows that people can be set in a promotion mode, where focus is on seeking to approach a desired state (to win), or in a prevention mode, focusing to avoid undesired states (not to lose). In promotion, people are eager to take risks increasing speed and decreasing accuracy. In prevention, strategic vigilance increases, decreasing speed and improving accuracy. When time and accuracy have to be compromised, one can ask which of these 2 strategies optimizes reward, leading to optimal performance. This is investigated here in a unique experimental environment. Decision making is studied in rapid-chess (180 s per game), in which the goal of a player is to mate the opponent in a finite amount of time or, alternatively, time-out of the opponent with sufficient material to mate. In different games, players face strong and weak opponents. It was observed that (a) players adopt a more conservative strategy when facing strong opponents, with slower and more accurate moves, and (b) this strategy is suboptimal: Players increase their winning likelihood against strong opponents using the policy they adopt when confronting opponents with similar strength.

Keywords: chess, decision making, response time, adaptation, speed–accuracy trade-off

Decision making involves the selection of a course of action among several alternative scenarios. An important aspect is whether the decision making mechanism is optimal, which can only be assessed in the context of a specific goal. For instance, a conservative policy in which decisions are only made when reaching maximal confidence in favor of an option may be optimal if the goal is to avoid errors. However, this policy is suboptimal if the goal is to maximize the number of correct responses (regardless of errors), where a less conservative policy that speeds the decision-making process increases the expected value (Green & Swets, 1966). Hence, the specific relevance of the outcomes of a decision, intended (maximizing correct choices) and unintended (making errors, wasting time, effort, or resources), depends on the specific context and goals in which decisions are made.

In this article, we investigate optimality of decision making in chess games played under intense temporal pressure. Rated chess games are played with a time control. Each player's clock starts with a specified time budget. While a player is deciding on their move, their clock time decreases, and the opponent clock stays still. A player who runs out of time automatically loses, unless the opposing player has insufficient material to mate, in which case the game is a draw. A timed-out player loses, even when holding a completely winning position with immediate mate to come. Hence, in chess under time control, the objective is to mate the opponent in a finite amount of time or, alternatively, time-out of the opponent with sufficient material to mate.[1] Different initial ranges of time budgets are given specific names: *classic* (about 2 hr), *rapid* (10 to 60 min per player), *blitz* (3 to 5 min per side) and *lightning* (when the time budget per side is less than 3 min).

In this study we concentrate on blitz games with a total time budget of 3 min per player without increment. This time control might be considered to be so fast that tactics and skill are secondary to quick moves. However, the prevalent view in chess exper-

Diego Fernandez Slezak, Computer Science Department, School of Sciences, University of Buenos Aires, Buenos Aires, Argentina; Mariano Sigman, Physics Department, School of Sciences, University of Buenos Aires, Buenos Aires, Argentina.

We thank Stockfish developers for the open source chess engine and for promptly responding to many questions, which have greatly contributed to the development of this work. We also thank the developers and administrators of FICS and FICSGAMES for continuously supporting a free chess server of the highest quality and for being completely open to experimentation. This work was supported by the Human Frontiers Science Program. The authors thank Jorge Moll, Agustín Gravano, Juliana Leone, and Cecilia Calero for fruitful discussions and useful comments on the article. The computing power was partially provided by Centro de Computación de Alto Rendimiento (CeCAR), http://pme84.exp.dc.uba.ar/ and IBM BlueGene/P.

Correspondence concerning this article should be addressed to Diego Fernandez Slezak, Departamento de Computación, Pabellón I, Ciudad Universitaria. (C1428EGA) Buenos Aires, Argentina. E-mail: dfslezak@dc.uba.ar

[1] Many tournaments are played with an "increment" that implies that in each move made, the player gains a fix number of seconds to its time budget. Under this rule, the total time is unbounded, but the total time per move remains bounded.

tise has ruled out this possibility. Expert players excel specifically at rapid object recognition abilities (Burns, 2004; Gobet & Simon, 1996b) and, hence, under temporal pressure are expected to further amplify the differences with weaker players. Time pressure also increases the possibility of blundering, even in grandmaster play (Chabris & Hearst, 2003). Indeed, 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 & Wagenmakers, 2005). 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, and, as a result, ratings at different time controls tend to be similar (Burns, 2004), indicating that, even when the time budget is severely limited, chess skill relates highly to the outcome. Hence, blitz constitutes an ideal context in which correct outcomes (good moves) and speed of choice have to be compromised to achieve the goal.

The specific aim of this work is to investigate (a) whether there is change in the decision-making policy in blitz when players face a stronger opponent and, (b) if so, whether this change is optimal assuming that the objective of the player is to win the game. To contextualize the current investigation in psychological theory, in the next sections, we review three relevant concepts.

Speed–Accuracy Trade-Off

In most psychological experiments, a relation is found between the time spent in a decision and its outcome known as the speed–accuracy trade-off (SAT; Woodworth, 1899). Participants in reaction-time (RT) tasks must adjust their performance to achieve an appropriate balance between speed and accuracy (Wickelgren, 1977).

Decision making has been widely explained by models in which evidence is progressively accumulated to a threshold (Luce, 1991). These models can easily explain the specific compromise between speed and accuracy in decisions based on noisy evidence. Raising the decision threshold diminishes the effect of noise in the outcome of the decision (it is formally equivalent to decreasing the effect of the noise). Naturally, raising the decision boundary also increases the time to reach the threshold resulting in slower response times (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Bogacz & Gurney, 2007; Bogacz, Wagenmakers, Forstmann, & Nieuwenhuis, 2010). This theory argues that fast and slow decisions are explained with a single model with a varying continuous parameter. But there are other potential alternatives to explain the SAT. For instance, according to the fast guess theory (Ollman, 1966; Yellott, 1971), in any individual decision a subject can emit either a random guess with short latency or a stimulus-controlled response at considerably longer latency. Under speed conditions, a subject increases the proportion of fast random guesses to stimulus-controlled responses. This theory also predicts the SAT relation, arguing that fast and slow decisions are of different kinds. However, several experiments have shown consistently that specific patterns of the SAT function cannot be accounted by different proportions of fast guesses (Pachella, 1974; Reed, 1973; Swensson, 1972).

Simple perceptual judgment has been investigated by Shadlen and collaborators. In these experiments, participants have to decide the direction of motion of a cloud of dots (Gold & Shadlen, 2007), where only a fraction of the dots move in a coherent direction and the remaining set of dots move randomly. As the proportion of random moving dots increases, the decision becomes more difficult. Under time pressure, participants can make these decisions quickly but at the cost of making more errors.

Gold and Shadlen (2002) tested a specific version of this task in which participants control the pace of response. The number of trials per unit of time increases as participants respond faster. Subjects receive a reward if the choice is correct, and there is no penalty for errors. Intuitively, the optimal threshold would establish a compromise between the two factors (speed and accuracy), limiting the reward rate. If the threshold is too low (relative to the difficulty of the task) the subject is not accurate, and if the threshold is too high, the subject is too slow. In this simplified version, the value of the decision threshold that maximizes the reward for a given fraction of coherent dots can be explicitly calculated (Bogacz et al., 2006), and it has been shown that participants rapidly and spontaneously adapt the threshold to its optimal value (Gold & Shadlen, 2002).

Although this suggests that in relatively simplified situations (a binary one-dimensional choice), the decision threshold achieved is optimal (in the context of maximizing the rate of reward), optimality is not found systematically in decision making. For instance, Bogacz, Hu, Holmes, and Cohen (2010) showed that in a simple modification of the random dot task, in which the time delay between trials is varied, subjects typically set their threshold at higher values than predicted by reward maximization, leading to more conservative policies. Moreover, Balci et al. (2011) trained people with multiple signal qualities and showed that it takes many sessions to train participants to behave in an optimal way, which shows that, in the more typical case, optimal threshold adaptation does not occur on the fly.

The beauty of the previous examples is that optimality can be determined explicitly by quantitative models of decision making. But these oversimplified situations are relatively far from real-life decision making, which incorporates desires, fears, social interactions, and core motivations idiosyncratic of human cognition (Kahneman & Tversky, 1979).

Kahneman and Tversky proposed an heuristic description of human decision making, with special emphasis in errors leading to suboptimal behavior which result from fear to take risks (referred as *risk aversion*). Their work is synthesized in *prospect theory*, proposing that the value function is steeper for losses than for gains, so that the subjective experience of pain from a loss is greater than the experience of pleasure from winning. In other words, that people tend to prefer sure gains over bets, even when the bet has a higher expected value than the sure gain (Tversky & Kahneman, 1981). Fear of losing leads to an overly cautious behavior, which turns out to be suboptimal.

Tversky and Kahneman also showed that the choices that one makes are affected by the way the problem is framed, leading to suboptimal behavior (Tversky & Kahneman, 1981). In their classic demonstration, they asked a group of participants to choose between two alternative programs to combat the disease: One had an exact value (200 people will be saved), and the other had a distribution of probabilities with risk involved ("there is a one-third probability that 600 people will be saved"). When the program was presented, as in the previous example, framed in terms of the number of people saved, participants preferred exact program. When the program was framed as the number of people who would die, participants instead chose the risk-taking decision. More generally, decisions that lead to gain are often risk averse, whereas choices involving losses are generally risk taking. This behavior can convincingly lead to suboptimality, when the bet has a higher expected value than the sure gain (Kahneman & Tversky, 1979; Tversky & Kahneman, 1981).

Regulatory Focus: Setting a Speed–Accuracy Trade-Off Policy

As reviewed in the previous section, prospect theory and the framing effect have convincingly demonstrated that human decision making can violate rational requirements of consistency, not maximizing the utility.

Prospect theory shows that decision makers weight gains and losses unevenly. Decision frame theory goes beyond this asymmetry, showing that decision makers can become suboptimal because of a specific formulation of the decision problem, norms, habits, and personal characteristics of the decision maker.

Recent psychological research has shown that a decisive factor determining the decision frame is the regulatory focus: the specific way in which someone approaches pleasure and avoids pain. Regulatory focus theory differentiates between two focus: a promotion-focus based on hopes and accomplishments (gains) and a prevention-focus based on safety and responsibilities (non-losses).

Cesario, Grant, and Higgins (2004) studied the interaction between regulatory focus and message framing. They primed either a promotion or prevention focus in participants. When a promotion focus was adopted, messages emphasizing the health benefits of exemplary fruit and vegetable consumption were more compelling. When a prevention focus was adopted, messages emphasizing the health complications of deficient fruit and vegetable consumption became more compelling. This is predicted by regulatory focus theory (Higgins, 1997), which proposes that regulatory focus differ in their strategic inclinations for attaining desired end states. In a promotion focus, the end state is desired, and the decision-maker goal is to approach it. In a prevention focus the end state is undesired, and the decision maker goal is to avoid it.

This idea seems abstract but becomes clear in the following example, which is essential in our study. Imagine a game played for money between two players. If Player A is in a promotion focus, he will play to win (achieving the desired state), which typically will involve taking risks and being aggressive. If Player A is in a prevention focus, he will play to avoid losing (avoiding the undesired state), which will reflect a cautious way of playing, avoiding errors that might lead to defeat.

Now consider the case in which both players have different strengths. The rules of the game are such that the weaker player plays at higher odds, receiving more money for a win than the stronger player. In case of a draw, the weaker player wins money (less than for winning), and the stronger player loses money. Should the weaker player shift to a more prevention or promotion focus (relative to his original bias of course)?

We reasoned that the weaker player might try not to lose and, hence, shift to a more prevention focus, because by simply not losing, he wins money. Conversely, the stronger player is forced to win if he does not want to lose money, which should orient him on a promotion focus. We refer to this hypothesis as the *draw value* shift of regulatory focus.

Conversely, players could shift in regulatory focus based on the decision stakes (i.e., the money earned by the weaker player if he wins, which is the same in value to the money lost by the stronger player if he loses), becoming more prevention oriented as the stakes grow higher. We refer to this hypothesis as the *stake value* shift of regulatory focus.

The two hypotheses have qualitatively different predictions in the shift of regulatory focus as a function of the difference in strength.

- The *draw value* hypothesis predicts a monotonic dependence of the shift in regulatory focus. When a player confronts a stronger opponent he should shift to prevention (trying to assure a draw). The greater the difference, the greater the shift should increase since the expected value of the draw is higher. Conversely, when a player confronts a weaker opponent he should shift the policy used (to avoid drawing), becoming progressively more promotion oriented as the difference with the weaker increases avoiding draws.
- The *Stake Value* hypothesis predicts a U-shape dependence of the shift in regulatory focus. When a player confronts an opponent of very different strength (either stronger or weaker), the stakes grow higher, and, hence, according to this hypothesis, the player should become more prevention oriented.

The experimental challenge to distinguish between these two hypotheses is to measure parametrically the shift in regulatory focus as a function of the change in strength. For this we bring together the two main ideas outlined in this review: the regulatory focus and the speed–accuracy trade-off. These two observables are intrinsically related: Promotion produces a strategic eagerness to achieve the end state, increasing speed and decreasing accuracy. Instead, prevention focus concerns avoiding undesired states, with a strategic vigilance, decreasing speed and improving accuracy (Förster, Higgins, & Bianco, 2003). Hence, the prediction of a U-shape or a linear function of shift in regulatory focus translates directly into the same prediction for RT and accuracy (see Figure 1).

In this study, we use chess games (played for points instead of money) to discriminate between these hypotheses and to test the optimality of shifts in regulatory focus. In the next section, we revise broadly how chess has been used to understand decision making and other aspects of the workings of the human mind.

Chess as Vehicle to Understand Cognition

Chess thinking has been a very informative path to the inner operation of effective thought. Adriaan De Groot's seminal book *Thought and Choice in Chess* (de Groot, 1965) is considered by many psychologists to be one of the greatest works on human thinking (T. Shallice, personal communication, December 2007).

Charness (1992) reviewed the impact of chess research on cognitive sciences in three different roles: (a) as a subject of inquiry in its own right (e.g., de Groot, 1965; Schultetus & Charness, 1999; Simon & Chase, 1973; van Harreveld, Wagenmakers, & van der Maas, 2007); (b) as a convenient environment for the study of complex cognitive processes, such as perception, problem solving, and memory (e.g., Chase & Simon, 1973; Gobet & Simon,

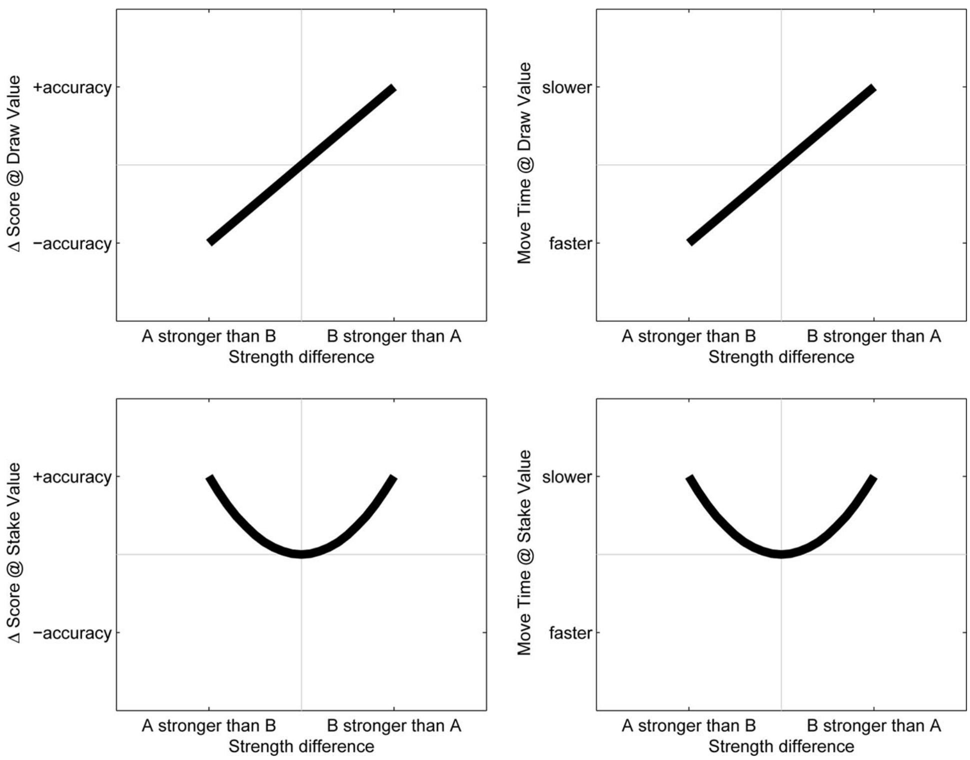

Figure 1. Hypothetical shifts of regulatory focus when players confront opponents of different strengths. Two hypotheses are proposed: *draw value* (top row), with promotion focus (less accurate and faster playing) when confronting weaker opponents and prevention focus (more accurate and slower playing) with stronger opponents, and *stake value* (bottom row) U-shaped dependence of the shift in regulatory focus, with shift to prevention (more accurate and slower) focus as stakes increase.

1996a; Reingold, Charness, Schultetus, & Stampe, 2001); and (c) as a convenient environment for exploring and developing theories about search mechanisms (e.g., Gobet & Simon, 1996b; Saariluoma, 1995).

In this article, we incorporate new dimensions in which chess may be useful in the cognitive sciences: decision making, optimality, and regulatory focus. In addition, our study has several important methodological differences from the vast majority of previous investigations using chess to understand the human mind. Our study relies on a vast corpus (30,000,000 games, about 2,400,000,000 decisions), capitalizing on a broad worldwide tendency of people to play chess and on the existence of servers that accumulate this data. Our study in this sense seeks statistical emergents of potentially subtle effects, which may be detectable only with a remarkable number of observations and might remain undetected in small sample sizes typical of laboratory studies. Our study also relies on the current vast capacity to analyze the value of chess positions with high accuracy and on powerful computing resources to evaluate thousands of millions of decisions. This results in a very unique experimental condition in which players are free to choose from among a virtually infinite tree of options, and nonetheless, the value of the position can be measured directly and does not need to rely on indirect measurements such as self-reported post-choice satisfaction, which have been often used to estimate the goodness of choice, as shown in Dijksterhuis, Bos, Nordgren, and Van Baaren (2006).

Another virtue of our chess-based decision making corpus is that the strength of each player is also well documented. Each registered user has associated a rating that indicates the chess skills strength of the player, represented by a number typically between 1,000 and 3,000 points, calculated using the Glicko rating system (Glickman, 1999). Briefly, the rating is a dynamic variable that is updated after each game played that represents the player's strength in a very confident manner. The update in rating constitutes a reward (players vigorously try to improve their ratings), which is, as all other variables, registered in our corpus.

In fact, reward in blitz games works exactly as in the money games described in the previous section, where rating points play the role of money. Players receive rating points (instead of money) in proportion to the difference in their strength. In the case of a draw, the weaker player wins points (less than for winning), and the stronger player loses them (Glickman, 1999).

We investigate the policies used by players when confronting stronger opponents (SO policy) and when confronting players of equal strength (ES policy). By investigating how RT and accuracy change in both policies, we can derive shifts in regulatory focus and, thus, verify which hypothesis—draw or stake value—governs players' behavior. We use the term *fearful* to refer to the policy adopted in prevention focus (avoiding risks, high decision threshold, slower, and more accurate) and *fearless* for the policy adopted by a player in promotion focus (risk taking, lower decision threshold, faster, and less accurate). Using these terms, our objective can be rephrased,, asking whether the SO policy is fearless or fearful. By a quantitative analysis of the winning likelihood of each configuration (a certain position value, a certain time for each player, and the rating of both players) we can also determine which of these policies is optimal (i.e., increases the expected win of rating points).

Method

All games were downloaded from FICS (Free Internet Chess Server; http://www.freechess.org/), a free ICS-compatible server for playing chess games. Site administrators offered free access to full database of games since 2008, with more than 300,000 registered users and more than 30 million games. This information was stored in a PostgreSQL database (http://www.postgresql.org/). The games include players involved, their ratings, moves, and times of each move in millisecond precision.

We added the evaluation of each move. 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 wins, 0 is the result is a draw, and –1 if black wins. An ideal evaluation function exists for other type of games, such as checkers, that are 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 Claude 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 one point for a pawn, three points for a knight or bishop, five points for a rook, and nine points for a queen. The king is sometimes given an arbitrary high value, such as 200 points (Shannon, 1950), or any other value that 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 that assigns a score (often also referred as value) to 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.

We used Stockfish (Romstad et al., 2011), an open source chess engine, to analyze the moves and calculate the score. The analysis consists of a finite tree-driven exploration of successive moves, up to a predefined depth of move number; we stored analysis with multiple depths, from six to 12 plies, into the database. 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 been played by both sides. According to the engine's exploration, other moves would worsen the score.

In this article, we used the score calculated by the maximum depth analyzed: 12 plies. The $\Delta Score$ was calculated as the difference between the given position score at depth 12 and the next ply at depth 11 to compare results with the same depth in the decision tree. Hence, the most accurate moves have $\Delta Score$ values close to zero, and very inaccurate moves correspond to high negative numbers (blunders). Because of the high performance of the computer engine, the fraction of moves with a positive value of $\Delta Score$—which provides an estimate of the goodness of the engine relative to the players of the database—is very low (Sigman, Etchemendy, Fernandez Slezak, & Cecchi, 2010).

The calculation of the score of a games may take between 1 and 3 min on a standard desktop computer. For the analysis used in this article, more than 1 million games were studied. If each game took 1 min to be analyzed, almost 2 full years would have been necessary to finish all of them. To accelerate this procedure, score calculation was parallelized using MPI and run in BlueGene/P and an off-the-shelf Beowulf cluster located at the University of Buenos Aires, Buenos Aires, Argentina.

We focused on players between 1,400 and 1,900 rating points. For all figures and statistical data reported in this article, we performed the same analysis on seven different sets each comprising 150,000 blitzes, and then results were averaged for these seven independent measures. Error bars indicate the standard error of these seven measures. All games considered in this article correspond to 180 s total time without increment.

During the analysis of the different variables in this article, we performed several filtering into the moves database. Filtering moves was performed to assure that all conditions were equal except the variable of interest (i.e., it is a way of selecting moves for which score, time left for white, time left for black. . . are matched and only the variable of interest changes). Because of the granularity of ratings, move times, and scores, selecting strict values results in almost single samples (if any) in the database. To have enough moves to be useful for statistics, we defined intervals for ratings, move times, and scores dimensions.

To illustrate this, we present a simple example. In Figure 4, we study the evolution of an average game. The moves used to calculate the move time where filtered first by players rating ([1,400–1,500]) and opponents rating ([1,600–1,700]). To match score and time-left, intervals were defined so that selected moves corresponded to those with score $S \pm \in_S$ and time-left $T \pm \in_T$ (i.e., $\in_S$ and $\in_T$ were, respectively, set to 0.2 and 5 s). These values assured sufficient moves for all conditions, but all the results reported here were robust to changes in these values.

Game trajectories (evolution of time and score throughout the course of the game) were calculated from moves played with SO and ES policies, following Algorithm 1 (see Figure 2).

In the SO policy, moves are obtained from games in which a player plays against a stronger opponent. For instance, in a representative game that opposes players rated 1,400 and 1,600, after 10 iterations (assuming that score is -1.5 and time-left is 150 s and 152 s for player and opponent, respectively), the *meanDuration* would be calculated as

$$\begin{aligned} meanDuration = \langle Duratinons | PlayerRating &= 1400 \\ \wedge\ MoveNumber = 11 \wedge Score &= -1.5 \\ \wedge\ TimeLeft &= 150\text{sec} \\ \wedge\ \underbrace{OpponentRating = 1600\rangle}_{SO\ policy} \end{aligned}$$

and *opponentMeanDuration* would be calculated as

Algorithm 1 Calculated ***ES*** game-play of the first 30 moves with player rating = pr and opponent player = $pr + \Delta rating$

```
T(0) = 180 sec
S(0) = 0
for i=1:30 do

  {Player's turn}
  meanDuration =
  ⟨Durations | PlayerRating = pr ∧ MoveNumber = i ∧ Score = S(i) ∧
        ∧ TimeLeft = T(i) ∧ OpponentRating = PlayerRating ⟩   [ES policy]
  meanDeltaScore =
  ⟨ΔScore | PlayerRating = pr ∧ MoveNumber = i ∧ Score = S(i) ∧
        ∧ TimeLeft = T(i) ∧ OpponentRating = PlayerRating ⟩   [ES policy]
  T(i) = T(i-1) - meanDuration
  S(i) = S(i-1) + meanDeltaScore

  {Opponent's turn: same procedure, but with SO policy, PlayerRating = pr + Δrating
  and OpponentRating = pr}

  ...

end for
return T,S
```

Figure 2. Algorithm 1. Calculated *ES* game-play of the first 30 moves with player rating = *pr* and opponent player = *pr* + Δ*rating*.

$$\begin{aligned} opponentmeanDuration = \langle Duratinons | PlayerRating &= 1600 \\ \wedge\ MoveNumber = 11 \wedge Score &= -1 \\ \wedge\ TimeLeft &= 152\text{sec} \\ \wedge\ \underbrace{OpponentRating = 1400\rangle}_{SO\ policy} \end{aligned}$$

On the other hand, a the same game-play using an ES policy would change only in the calculation of players averages; for example, the *mean Duration*:

$$\begin{aligned} meanDuration = \langle Duratinons | PlayerRating &= 1400 \\ \wedge\ MoveNumber = 11 \wedge Score &= -1.5 \\ \wedge\ TimeLeft &= 150\text{sec} \\ \wedge\ \underbrace{OpponentRating = 1400\rangle}_{SO\ policy} \end{aligned}$$

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 by browsing them as a guest. That is, the data are already anonymized. Moreover, in our storing process, we further anonymized the data by stripping all information except the player's ranking. Individual consent was therefore not sought because of this double layer of anonymity, along with the public, open nature of the website.

Results

Changes in SAT When Opposing Players of Different Strength

Our aim is to determine whether regulatory focus shifts toward a more prevention or promotion focus when players face opponents of different strength. This can only be measured indirectly, calculating shifts in the speed–accuracy trade-off.

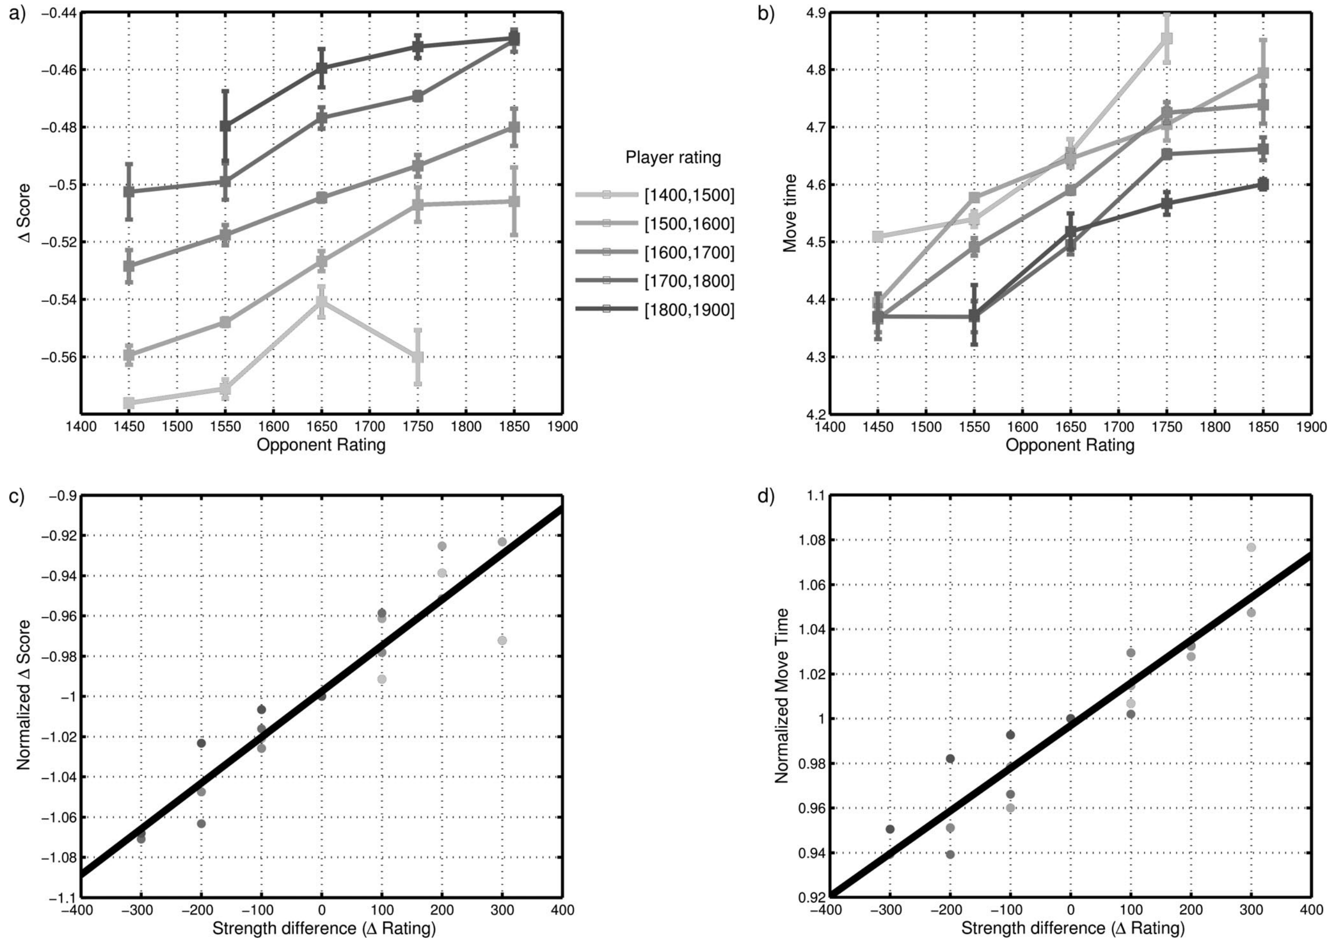

Figure 3. Shifts in response time (move-time) and accuracy (Δ*Score*) when players confront opponents of different strengths. Average Δ*Score* (a) and move-time (b) were calculated for players between 1,400 and 1,900 rating points, confronting opponents between 1,400 and 1,900 rating points, both grouped in bins of 100 rating units: [1,400–1,500], [1,500–1,600], . . . , and [1,800–1,900]. Only move numbers between 10 and 30 were taken into account from 150,000 blitz games (3 min without increment). Spontaneous change in policy is observed: As the strength of the opponent increases, players spend more time making each move and make more accurate moves, consistent with the *draw value* hypothesis. In panels c and d, respective normalized accuracy and response-time are shown versus strength difference, confirming the draw value hypothesis.

To investigate shifts in SAT, we measured whether response times and accuracy (change in Δ*Score*) varied when players confronted opponents of different strength. We emphasize that Δ*Score* is a measure of the quality of a single, one-player move, independent of the opponent response. It is obvious that the winning likelihood will diminish when playing stronger opponents, since the opponent move will be progressively better. But here we do not take this factor into account. Instead, we only focus on one-player move and investigate if, when presented with the same position, they play differently as a function of the opponent strength.

We selected moves between 10 and 30 from 150,000 blitz games (3 min without increment) with players between 1,400 and 1,900 rating points. We split the moves, taking into account the player ratings, into five different sets: [1,400–1,500], [1,500–1,600], . . . , and [1,800–1,900] rating points (see Figure 3, data series in color lines). For each data series, we calculated the average move-time and Δ*Score* versus the opponent's rating, again in bins of 100 rating points.

Analysis of performance revealed a highly expected main effect: Strong players play more accurately and faster than weak players (see Figure 3a). Accuracy is reflected by an increase of score with player rating. An equivalent way of expressing this result would be to set a blunder threshold (typically a move is considered a blunder when the loss of score after the move is more than 2, but this threshold is arbitrary). As the average score increases, the likelihood of blundering decreases, as shown in Sigman et al. (2010), which implies that, as expected, stronger players have less probability of blundering.

Interestingly when a player confronts a stronger opponent accuracy increases. To quantify this observation—capitalizing on the vast amount of the data—we compared pairs of decisions for which all variables were identical (time left to the player, time left to the opponent, evaluation of the position, player's rating) and only differed in the rating of the opponent, which could either be the same, 100 or 200 rating points of difference.

We performed the same analysis for move duration, investigating how it varies with opponent strength when all other state variables of the game are matched. We found that, as observed for accuracy, duration increases in a highly significant manner as the difference between the player and the opponent rating increases (see Figure 3b). For instance, a player with 1,650 rating points during the middle game (between moves 10 and 30) takes (on average) less than 4.4 s per ply when facing a weaker opponent and more than 4.7 s when facing stronger opponents.

To assess which hypothesis was adopted by players (Draw or Stake value), we calculated the normalized accuracy—$\Delta Score^{N} = \Delta Score[\Delta Rating)/\Delta Score(\Delta Rating = 0)]$, where $\Delta Rating$ = opponent's rating minus player's rating—versus strength difference with opponent (see Figure 3c). We found that, as predicted by the Draw value hypothesis, accuracy increases linearly with opponent's rating increase. Analogously, we calculated the normalized response times (RT^{N}) versus strength difference. Again, consistent with the Draw value hypothesis, we found that response time increases linearly as opponent's strength increases (see Figure 3d).

We calculated the linear regression of change in $\Delta Score^{N}$ and RT^{N} versus difference with opponent's strength. Figure 3c shows $\Delta Score^{N}$ (color points) and the linear regression (black line, $R^2 = .88$, $p < 10^{-12}$). Change in response time (RT^{N}) is shown in Figure 4d (color points) and its linear regression (black line, $R^2 = .89$, $p < 10^{-12}$).

Our results show that in rapid chess players spontaneously change their policy: As the strength of the opponent increases, they become more conservative in the speed–accuracy axis (i.e., they spend more time making each move and make more accurate moves). This result denies the *stake value* shift of regulatory focus, since it would predict a consequent U-shape dependence in RT and accuracy. It is consistent with the *draw value* hypothesis. Regulatory focus shifts toward a more prevention oriented strategy as opponent's strength is higher as reflected by the adoption of a fearful policy, where decisions are slower and more accurate.

Optimality of Policy Shifts When Confronting Stronger Opponents

Since rapid chess is played with finite time, and since we have sufficient games in the database for each trio of (value of the position, time left for player one, time left for player two), it is possible to determine empirically whether this switch in policy is optimal (Sigman et al., 2010). Figure 4 gives a hint indicating that this shift may be suboptimal: Although the amount in time increases steadily with opponent strength, accuracy shows a plateau, suggesting that the cost in time may not compensate the loss of time.

To quantitatively determine if policy adaptation is optimal, we investigated the winning likelihood in simulated game trajectories (evolution of time and score throughout the course of the game) played with SO (henceforth referred as *fearful*) and ES (henceforth referred as *fearless*) policies.

Given two ratings (player and opponent), we can calculate the mean move time and $\Delta Score$ for each game status (time-left of each player and score). From an initial state of the game, we can simulate a game trajectory by iterating representative moves of each player under different policies.

Specifically, the trajectory of a representative game is calculated iteratively as follows: *Time-left* is initially set to 180 s (all games analyzed here were 3 min games without increment), and *score* is set to 0. A move is made with a duration D and with a change in score $\Delta Score$, considering the mean of durations (respectively, change of score) for all moves in the database with the corresponding time-left T and score S. Iteratively, time-left and score evolved following the average duration and $\Delta Score$ (see Algorithm 1 for an example of calculated game-play using the fearless policy).

To create *fearful* game trajectories, we can just fix the players ratings (e.g., wr for white-rating; *br* for black-rating) and initial status and concatenate mean moves calculated from games with white and black players of ratings *wr* and *br*, respectively. Instead, we can simulate *fearless* policy game trajectories, by calculating the mean move from other game set, with players with the same rating (e.g., white and black ratings equal to wr). If white player was weaker than black player, from Figure 4 we know that, in the *fearless* policy, the white player would play faster and lose a tiny additional amount of value in each move.

To make this aspect of our experiment clear, we asked whether 1,500 rated players play optimally against 1,700 rated players. The key aspect of this experiment is that, at any moment of the game (e.g., it is white's turn, with 60 s on the clock; the opponent has 80 s, and the position is slightly favorable with a score of 0.5), we have sufficient data to compute precisely the mean time taken and the average quality of this move. This is done for two different distributions of moves: considering only moves belonging to games in which 1,500 rated players confront 1,700 players (fearful policy) or only moves belonging to games in which 1,500 players confront opponents of the same strength (fearless policy). The algorithm then iterates time and score according to the measured average values. Critically, the response of the opponent does not depend on the player policy. The response of the 1,700 player will be computed from the moves belonging to all games in which a 1,700 player confronts a 1,500 player after filtering for the corresponding time-left of both players and score. The explicit algorithm is included in the Method section.

Figure 4 shows the evolution of the game according to both policies for players with 1,450 ± 50 rating points facing players with 1,650 ± 50 rating points. The fearful policy represents the evolution of an average game of the database. Time budget decreases monotonically, and after 30 moves, players have used, on average, about half of their time budget. Note that the time-left curve becomes steeper, indicating that more time is spent in the middle game than in the opening stage of the game, as shown in Sigman et al. (2010). The score also decreases throughout the course of the game, which simply indicates that a player is facing a stronger opponent, and as the game progresses, his position worsens. The evolution of a game with a fearless policy follows basically the same trend, with consistent departures (error bars are plotted on the time-left curves, but they are so small that they are almost invisible). When, playing according to the fearful policy, the time budget decreases more rapidly, reaching a difference of about 3 s after 30 moves. Consistently, when playing with the fearful policy, the score remains better than when playing with the fearless policy, reflecting the integration over moves of a weak but significant difference in accuracy in each move of the sequence.

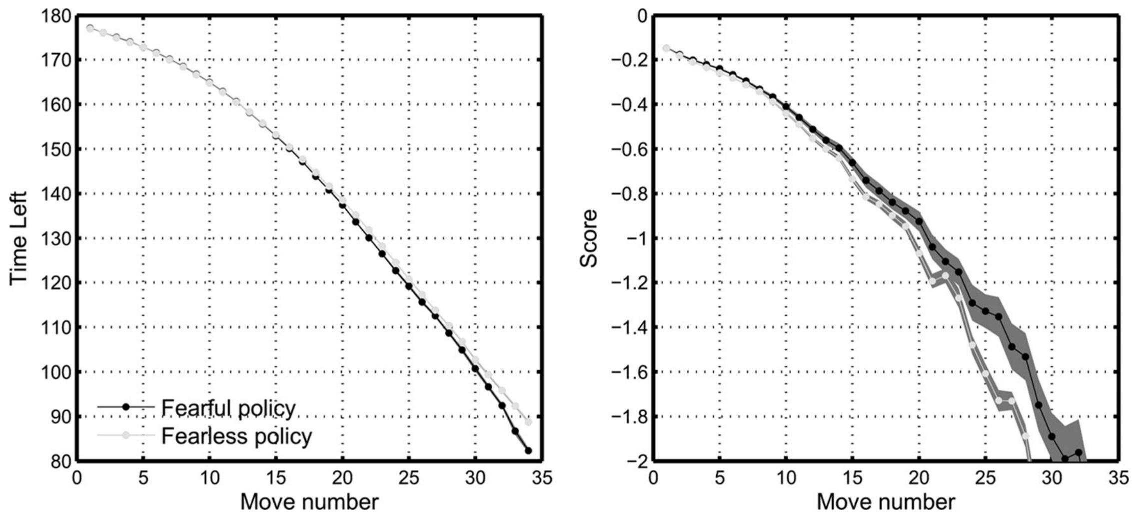

Figure 4. Simulated game trajectories (evolution of time and score throughout the course of the game) played with *fearful* and *fearless* policies, calculated for players rated between 1,400 and 1,500 against opponents with rating 1,600–1,700. In black dots, *fearful* policy represents the evolution of an average game of the database. In gray dots, players' *fearless* policy calculated from games confronting opponents of the same strength: [1400;1500]. Shadow shows the error bars on both panels; errors are so small that, in left panel, they are almost invisible. After 30 moves, a player following a fearless policy stands with a slightly worse position but with more time available.

In summary, after 30 moves, a player following a fearful policy (comparatively slower and more accurate moves), on average, stands with a position that is slightly better than if he or she would have played with a fearless policy. But this does not come without a cost; he or she also has, on average, less time remaining for the rest of the game. Are his or her winning chances better or worse?

To determine the efficiency of both policies, we compute the average result as a function of the time-left to each player, the score, and the rating of both players. How all these factors combine to form an average result can be determined empirically from the database. This is done by simply computing, on each move number, the average result of all games that agree with these factors. It is important to emphasize that, whereas in the fearless policy, the moves considered to evolve the game are taken from players of the same strength, the average result is calculated, for both policies, from the games in which the player confronts the stronger opponent.

We estimated the average result according to play dictated by both policies for seven independent experiments (each comprising 150,000 independent games) and for three different Δ*Rating* values: 100, 200, and 300 rating points (see Figure 5). For each game, the result has a value of 1 for a win, 0 for a draw, and –1 for a loss. Since we study games in which players face a stronger opponent, average results are negative, and, as expected, the average result decreases as Δ*Rating* growths.

For the 15 cases studied here (five different ratings and three different Δ*Ratings*), we observed that average result was greater for the fearless than for the fearful policy (note that all points in Figure 5 are above the diagonal). This observation is confirmed by an analysis of variance (ANOVA) with player rating ([1,400

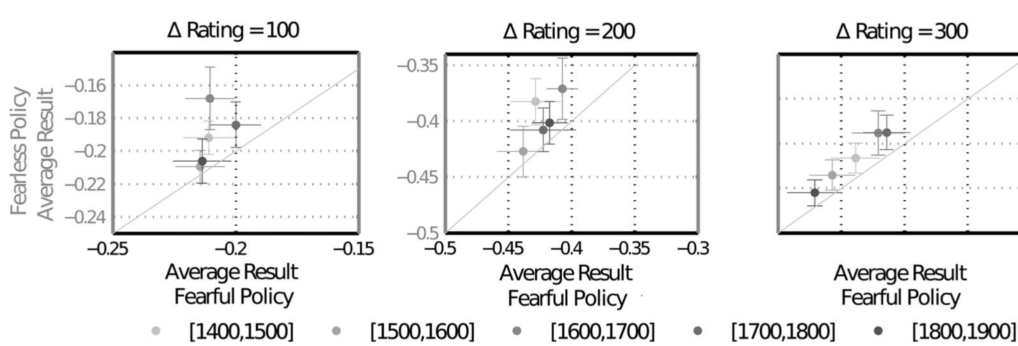

Figure 5. Winning likelihood in simulated game trajectories based on games played using fearless versus fearful policies. Trajectories were calculated for players between 1,400 and 1,900 rating points (grouped in bins of 100 rating points), confronting opponents with three different Δ*rating* values: 100, 200, and 300 rating points. All rating ranges show better estimated results using the fearless policy.

1,500], [1,500 1,600], [1,600 1,700], [1,700 1,800], [1,800 1900]) Δ*Rating* (100, 200, 300), and policy (fearful and fearless) as main factors. This shows that players confronting stronger opponents play slower and more accurately and that these factors combine in a suboptimal way.

The effect was extremely consistent, as for all ratings and Δ*Ratings*, we observed that winning chances increased with the fearless policy (see Figure 5). This effect was confirmed by a *t* test comparing the average result of the fearless and fearful policies for all ratings and Δ*Ratings*, as in Figure 5, $t(14) = 4.18$, $p < .001$. Although this effect is extremely consistent (winning likelihood increased for every single rating and Δ*Rating*), it must be emphasized that the effect size is relatively small and is only observable after averaging a significant number of games played. However, the majority of these players play a great number of games, suggesting that a concrete intervention based on these observations may have a concrete benefit on weaker players' performance.

Discussion

We proposed that players adapt their regulatory focus, i.e., the balance between pursuing positive outcomes—a win—and avoiding negative ones—a loss. Argument on how regulatory focus might be set lead to two different hypothesis: (a) the *draw value* hypothesis, according to which a player confronting a stronger opponent should shift to prevention (trying to avoid losing to assure a draw), and (b) the *stake value* hypothesis, according to which a player confronting an opponent of very different strength (either stronger or weaker) should become more prevention oriented, because the stakes grow higher.

We showed that players of different levels adapt their speed–accuracy policy (i.e., their regulatory focus) when confronting opponents of varying strengths. When facing strong opponents, players shift toward a prevention focus, which results in a fearful decision-making policy, playing slightly slower and more accurately. This result denies the *stake value* hypothesis and provides evidence in favor of the *draw value*. As players face stronger opponents, they concentrate more on not losing, being more cautious, and playing slower and more accurately.

If players play more accurately, why does winning likelihood decrease? This can only be understood in terms of the rules of the game. The opponent has to be mated in a relatively short (180-s) time budget. Hence, optimizing winning chances requires a balance between playing good and playing fast and our results show that, empirically, the prevention focus does not maximize the gain of these competing factors. The price of the time spent is too high for the corresponding gain in value. It is clear that this may only be valid in games played under strong time pressure (or, as often happens, in games that initiate with a generous time budget but that become a time-pressure game after a player "sleeps on his clock"). As emphasized in the introduction, optimality can only be argued with a specific goal and context. Here we show that, with strong time pressure, becoming prevention oriented is, in fact, suboptimal. Our results show that the extra time spent is too costly for the improvement in the position to which it leads. The possibility of measuring the *price of a second* is unique in this experimental setup and turns out to be a very difficult question to approach in general decision-making scenarios. We could capitalize on the experimental setup of rapid chess that combines virtually infinite data (thousands of millions of decision), time pressure, and very refined algorithms to compute the accuracy of a decision (of the move made) and its corresponding change in value.

A specific analysis of the game played under different policies shows that suboptimality emerges when both players have more than 1 min remaining in their clocks, before time pressure converts the game to what is essentially a moving piece lottery. Hence, our results do not rely on decisive outcomes from this extreme time-pressure situation, in which the most reasonable strategy is simply to make any legal move to avoid being flagged and to try to flag the opponent. There are also some exceptions to this, with impressive demonstrations of how players can master a mating line with only a couple of seconds remaining on the clock. In particular, most chess clients used for online chess have an option used by virtually all players by which a player can make a *premove* before the opponent has played. If the premove is legal it is made without loss of time.

Generally, time pressure tends to impose a more intuitive (impulsive) strategy, rather than thoughtful decisions made with plenty of time (Kim & Lee, 2011), based on a tendency to seek an inferior but more immediate reward. Impulsive actions can be produced when too much emphasis is placed on speed, rather than accuracy, in a wide range of behaviors, including perceptual decision making. In our experiment, under strong time pressure, players adopt longer deliberating decisions when facing strong opponents, becoming less impulsive. Since time primes here, against what intuition may dictate, this lack of impulsivity turns out to be suboptimal.

Our results also rule out other factors, which have previously been considered instrumental in determining the suboptimal character of decisions with deliberation (e.g., Dijksterhuis et al., 2006; Wilson & Schooler, 1991). When the space of choices involves too many factors (e.g., choosing a car in which a complex function of speed, size, prize, comfort, economy, etc. has to be maximized) the decision without deliberation is better, and certainly faster, than a decision based on rational thought (Dijksterhuis et al., 2006; Wilson & Schooler, 1991). We do not observe this dependence, even if chess presents without doubt a high-dimensional search space. We showed that when players face weaker opponents, they play faster and their winning chances increase. But this does not result from the fact that fast decisions without deliberation are better than slow, calculated rational choices. This is consistent with recent results showing that, 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 & Wagenmakers, 2005). On the contrary, when players play faster, the result of each move is worse on average. The only reason why this strategy is nevertheless optimal is, as mentioned earlier, because the time saved plays a more important role.

We have shown a change in policy when confronting stronger opponents and argued that this is an intentional (not necessarily conscious) shift associated with an adoption of a prevention focus. However, one may question whether there may be other possible causes for why players play more slowly and more accurately when confronting stronger opponents. A possible concern is that the adoption of the conservative policy may not necessarily be the result of the weaker player's individual decision but, rather, an emergent property of the dynamics of the game. Intuitively speaking, one may have no other choice but to fend off the blows, regardless of emotional state or inclination. This is a very important aspect that requires specific clarification.

After a few moves, a player is likely to have a worse position against a stronger opponent and a better position against a weaker opponent. Won positions might be easier to play (although this is not necessarily true), and hence, it may be thought that this may explain our results. However, our analysis takes this into account. When comparing the properties of a move time taken and change in score), we have matched all properties (time left for both players, the score of the position) and only varied, parametrically, the rating of the opponent. This is in part possible because of the vast amount of data in which we still have sufficient games in which a player plays a stronger opponent (at move *N*) with a score and time in each clock comparable to the expected value at move *N* when playing an opponent of the same strength. In spite of this effort to match all circumstances except rating, some issues may still remain. First, a game in which a weak player is holding an even game against a strong opponent is an atypical game in which he is playing above its chances. It may then be the case (and this is a possibility we cannot discard) that the shift in policy is enhanced when a weak player has made it far in the game without losing, and the desire of to not lose the game is exacerbated. A second argument is that even if the scores are equal, the stronger player may pose more difficult problems. This raises the issue of the complexity of a chess position (psychological, not computational, complexity), which is extremely hard to measure algorithmically. This was actually one of the original goals of the seminal work of de Groot (1965) and still remains a challenge for computer chess and artificial intelligence. Since psychological complexity is virtually impossible to measure precisely (as opposed to the score, which is very accurate) we could not match it in this study. But it could be possible that players take more time against strong opponents simply because the position they are presented with is harder to handle. However, an aspect of our data makes this hypothesis unlikely. If this was the case, we would expect that, when playing stronger opponents, players would take more time and also play worse. Or, at least, not be able to find a balance between time and accuracy in which more time results in proportionally better choices. The logic of the SAT compromise is evident when complexity is matched. On the contrary, when the complexity of a problem increases with a parameter—for instance, the distance of a numerical comparison (Dehaene, 1992)—as the decision becomes more difficult, subjects take more time and make more errors.

The finding that players shift strategy when facing strong opponents builds up in consonance with previous results by Förster et al. (2003) demonstrating that decisions appear to be influenced by the strategic inclinations of participants varying in regulatory focus rather than by a built-in trade-off. The observed suboptimality is reminiscent of the inability of expected utility to explain decision making under risk, as discussed by Kahneman and Tversky (1979). More generally, these results are in line with a widely demonstrated influence of emotions on decision making (e.g., Damasio, 2000; Isen & Patrick, 1983; Loewenstein, Weber, Hsee, & Welch, 2001). A corollary of our observations is the non-Markovian character of decision making during game play. In "perfect" play, at any given position, there is a "best move" that is independent of all factors beyond the specific position of the board, such as the precedent sequence of moves. In previous results, Sigman et al. (2010) showed non-Markovian violations in chess playing. The present observations extend these results. Although we do not have recordings of participants' awareness of such strategic change, it is very likely that it remains spontaneous and unconscious, probably reflecting heuristics, as proposed by Kahneman, Slovic, and Tversky (1982) and Ariely and Jones (2008), i.e., an estimate of how players should regulate their own play from an estimate of the opponent strength.

Practical (imperfect) play is actually not Markovian. For instance, backgammon software computes winning chances to decide cube policy, whether to double stakes or not, combining the value of the position (the Markovian component) and an estimation of a player strength (the non-Markovian component, since this estimation is based on the history of the game). Estimating opponent strength in chess is straightforward, since it is indexed by the opponent's rating. Using this information as an heuristic to combine and guide action could be an optimal non-Markovian way of playing. As it happens, however, the fear and respect that induce conservative play are—at least in fast chess, where time is too costly—not the right way to go.

References

Ariely, D., & Jones, S. (2008). *Predictably irrational: The hidden forces that shape our decisions.* New York, NY: Harper.

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision making criteria: Reward rate ultimately beats accuracy. *Attention, Perception, & Psychophysics,* 1–18.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review, 113,* 700. doi:10.1037/0033-295X.113.4.700

Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation, 19,* 442–477. doi:10.1162/neco.2007.19.2.442

Bogacz, R., Hu, P., Holmes, P., & Cohen, J. (2010). Do humans produce the speed–accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology, 63,* 863–891. doi:10.1080/17470210903091643

Bogacz, R., Wagenmakers, E., Forstmann, B., & Nieuwenhuis, S. (2010). The neural basis of the speed–accuracy tradeoff. *Trends in Neurosciences, 33,* 10–16.

Burns, B. (2004). The effects of speed on skilled chess performance. *Psychological Science, 15,* 442–447.

Cesario, J., Grant, H., & Higgins, E. (2004). Regulatory fit and persuasion: Transfer from "feeling right." *Journal of Personality and Social Psychology, 86,* 388.

Chabris, C., & Hearst, E. (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. doi:10.1207/s15516709cog2704_3

Charness, N. (1992). The impact of chess research on cognitive science. *Psychological Research, 54,* 4–9. doi:10.1007/BF01359217

Chase, W., & Simon, H. (1973). Perception in chess* 1. *Cognitive Psychology, 4,* 55–81. doi:10.1016/0010-0285(73)90004-2

Damasio, A. (2000). *The feeling of what happens: Body and emotion in the making of consciousness.* London, England: Vintage.

de Groot, A. (1965). *Thought and choice in chess.* The Hague, the Netherlands: Mouton.

Dehaene, S. (1992). Varieties of numerical abilities. *Cognition, 44,* 1–42. doi:10.1016/0010-0277(92)90049-N

Dijksterhuis, A., Bos, M., Nordgren, L., & Van Baaren, R. (2006). On making the right choice: The deliberation-without-attention effect. *Science, 311,* 1005. doi:10.1126/science.1121629

Förster, J., Higgins, E., & Bianco, A. (2003). Speed/accuracy decisions in task performance: Built-in trade-off or separate strategic concerns? *Organizational Behavior and Human Decision Processes, 90,* 148–164. doi:10.1016/S0749-5978(02)00509-5

Glickman, M. (1999). Parameter estimation in large dynamic paired comparison experiments. *Journal of the Royal Statistical Society Series C: Applied Statistics, 48,* 377–394. doi:10.1111/1467-9876.00159

Gobet, F., & Simon, H. (1996a). Recall of random and distorted chess positions: Implications for the theory of expertise. *Memory & Cognition, 24,* 493. doi:10.3758/BF03200937

Gobet, F., & Simon, H. (1996b). The roles of recognition processes and look-ahead search in time-constrained expert problem solving: Evidence from grand-master-level chess. *Psychological Science, 7,* 52–55. doi:10.1111/j.1467-9280.1996.tb00666.x

Gold, J., & Shadlen, M. (2002). Banburismus and the brain decoding the relationship between sensory stimuli, decisions, and reward. *Neuron, 36,* 299–308. doi:10.1016/S0896-6273(02)00971-6

Gold, J., & Shadlen, M. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30,* 535–574. doi:10.1146/annurev.neuro.29.051605.113038

Green, D., & Swets, J. (1966). *Signal detection theory and psychophysics.* New York, NY: Wiley.

Higgins, E. (1997). Beyond pleasure and pain. *American Psychologist, 52,* 1280–1300. doi:10.1037/0003-066X.52.12.1280

Isen, A., & Patrick, R. (1983). The effect of positive feelings on risk taking: When the chips are down. *Organizational Behavior and Human Performance, 31,* 194–202. doi:10.1016/0030-5073(83)90120-4

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge, England: Cambridge University Press.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica,* 263–291.

Kim, S., & Lee, D. (2011). Prefrontal cortex and impulsive decision making. *Biological Psychiatry, 69,* 1140–1146. doi:10.1016/j.biopsych.2010.07.005

Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin, 127,* 267–286. doi:10.1037/0033-2909.127.2.267

Luce, R. (1991). *Response times: Their role in inferring elementary mental organization.* New York, NY: Oxford University Press.

Ollman, R. (1966). Fast guesses in choice reaction time. *Psychonomic Science, 6,* 155–156.

Pachella, R. G. (1974). The interpretation of reaction time in information processing research. In B. Kantowitz (Ed.), *Human information processing* (pp. 41–82). Oxford, England: Erlbaum.

Reed, A. (1973). Speed–accuracy trade-off in recognition memory. *Science, 181,* 574. doi:10.1126/science.181.4099.574

Reingold, E., Charness, N., Schultetus, R., & Stampe, D. (2001). Perceptual automaticity in expert chess players: Parallel encoding of chess relations. *Psychonomic Bulletin & Review, 8,* 504. doi:10.3758/BF03196185

Romstad, T., Costalba, M., Kiiski, J., Yang, D., Spitaleri, S., & Ablett, J. (2011). *Stockfish, open source chess engine.* Retrieved from http://www.stockfishchess.com/

Saariluoma, P. (1995). *Chess players' thinking: A cognitive psychological approach.* New York, NY: Routledge.

Schaeffer, J., Burch, N., Bjornsson, Y., Kishimoto, A., Muller, M., Lake, R., . . . Sutphen, S. (2007). Checkers is solved. *Science, 317,* 1518. doi:10.1126/science.1144079

Schultetus, R., & Charness, N. (1999). Recall or evaluation of chess positions revisited: The relationship between memory and evaluation in chess skill. *The American Journal of Psychology, 112,* 555–569. doi:10.2307/1423650

Shannon, C. (1950). Programming a computer for playing chess. *Philosophical Magazine, 41,* 256–275.

Sigman, M., Etchemendy, P., Fernandez Slezak, D., & Cecchi, G. (2010). Response time distributions in rapid chess: A large-scale decision making experiment. *Frontiers in Neuroscience, 5,* 12.

Simon, H., & Chase, W. (1973). Skill in chess. *American Scientist, 61,* 394–403.

Swensson, R. (1972). The elusive tradeoff: Speed vs. accuracy in visual discrimination tasks. *Attention, Perception, & Psychophysics, 12,* 16–32. doi:10.3758/BF03212837

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211,* 453. doi:10.1126/science.7455683

Van Der Maas, H., & Wagenmakers, E. (2005). A psychometric analysis of chess expertise. *The American Journal of Psychology, 118,* 29–60.

van Harreveld, F., Wagenmakers, E., & van der Maas, H. (2007). The effects of time pressure on chess skill: An investigation into fast and slow processes underlying expert performance. *Psychological Research, 71,* 591–597. doi:10.1007/s00426-006-0076-0

Wickelgren, W. A. (1977). Speed–accuracy tradeoff and information processing dynamics. *Acta Psychologica, 41,* 67–85. doi:10.1016/0001-6918(77)90012-9

Wilson, T., & Schooler, J. (1991). Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology, 60,* 181. doi:10.1037/0022-3514.60.2.181

Woodworth, R. (1899). Accuracy of voluntary movement. *Journal of Nervous and Mental Disease, 26,* 743.

Yellott, J. (1971). Correction for fast guessing and the speed–accuracy tradeoff in choice reaction time* 1. *Journal of Mathematical Psychology, 8,* 159–199. doi:10.1016/0022-2496(71)90011-3

Received March 24, 2011
Revision received September 1, 2011
Accepted September 1, 2011 ■