UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Risk-taking in adversarial games: What can 1 billion online chess games tell us?

Permalink
https://escholarship.org/uc/item/403764rd

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN
1069-7977

Authors
Holdaway, Cameron
Vul, Ed

Publication Date
2021

Peer reviewed

Risk-taking in adversarial games:
What can 1 billion online chess games tell us?

Cameron Holdaway
Department of Psychology
UC San Diego
choldawa@ucsd.edu

Edward Vul
Department of Psychology
UC San Diego
evul@ucsd.edu

Abstract

Humans are social beings, and most of our decisions are influenced by considerations of how others will respond. Whether in poker or political negotiations, the riskiness of a decision is often determined by the variance of the other party's possible responses. Such socially-contingent decisions can be framed in terms of adversarial games, which differ from other risky situations such as lotteries because the risk arises from uncertainty about the opponent's decisions, and not some independent stochasticity in the world. We use chess as a lens through which we can study human risk-taking behavior in adversarial decision making. We develop a novel algorithm for calculating the riskiness of each move in a chess game, and apply it to data from over 1 billion online chess games. We find that players not only exhibit state-dependent risk preferences, but also change their risk-taking strategy depending on their opponent, and that this effect differs in experts and novices.

Keywords: risk taking; adversarial games; chess

Introduction

Humans are social beings, who spend much of their time making decisions based on how they think others will respond. Political and business dealings often involve this adversarial, game-theoretic flavor, where the quality of a decision depends on the other side's response. Studying risk in these adversarial games is interesting because the measure of risk is explicitly defined in terms of possible opponent responses. This can be contrasted with taking risks in a lottery game, in which the riskiness of a decision is based on some irreducible stochasticity in the world, and not the response of another person. While there are many studies of risk-taking behavior, which factors influence risk preferences, and how and when we choose to take risks, the adversarial aspect of risk-taking remains relatively unexplored. There is reason to suspect that risk preferences will differ when risk is arising from a sequentially-choosing adversary, because an adversary's choice seems more controllable. It seems like we might trick an opponent, and thus out maneuver them. Indeed, many biases evident in lottery choices (such as the gambler's or the hot-hand illusion) seem consistent with a mistaken belief that a sequentially-independent mechanical random process is somewhat sequentially-dependent, and perhaps non-stationary, as though it were generated more organically. Such biases suggest that even when playing against an immovable randomness generator, people believe that it might be swayed by prior outcomes in the sequence, like humans are (Brockbank & Vul, 2020). Thus, it seems all the more likely that people may treat risk generated from an unknown adversary — one that truly is organic — differently from a mechanical lottery process.

Risk-taking behavior and risk preferences have been widely studied in a variety of fields and it has been shown that taking risks goes beyond simply evaluating the probabilities and expected values (Slovic, 2000; Kahneman & Tversky, 2013). How one evaluates risks and chooses to accept them can depend on many environmental and individual factors, such as: expertise, current performance, and past performance. Experts and lay-people not only perceive risks differently (Bostrom, 1997), but may also behave differently in risky situations. For example, expert helicopter pilots exhibit a tendency for increased risk-taking in simulations of adverse events over less experienced candidate pilots (Thomson, Önkal, Avcioğlu, & Goodwin, 2004). Similar changes in risk-taking can be seen within the same person depending on circumstances and past performance. Tversky and Kahneman (1991) propose a reference-dependent model of decision making in which choices depend on a reference point and the status quo. Similarly, prior gains or losses can affect risk preferences, generating effects such as "playing with house money", wherein recent gains can lead to increased risk taking (Thaler & Johnson, 1990). Though these effects have been theorized or shown to hold in experiments where risk is determined by the environment, there remain open questions as to whether risk-taking behavior may be different when the risk arises from the choices of another person.

In the laboratory setting, lotteries, in which subjects choose between guaranteed and non-guaranteed payoffs, have been by far the most popular procedure for studying risk preferences (Kahneman & Tversky, 2013; Mata, Frey, Richter, Schupp, & Hertwig, 2018). While these lotteries come in a variety of different forms (for a comprehensive review of the many methods of this form see Harrison, Rutström, et al. (2008). The crucial component is that subjects choose between a risky option with high variance of outcomes and a safer option, and that the probabilities of payouts are generated from the world. These studies often have subjects choose between payouts with fully-specified probabilities (Harrison et al., 2008), or be required to infer the payouts and probabilities across repeated trials, as is the case in bandit tasks (Meyer & Shi, 1995), or other decisions from experience (Hertwig & Erev, 2009). The key difference between these settings and

many human-to-human interactions is that the source of risk comes entirely from stochasticity in the environment. In other words, the participant's choice is not dependent on possible responses of another agent.

Behavioral game theory, however, does seek to study these types of human-human interactions in which decision makers interact, and uncertainty arises from the indeterminacy of an opponent's choice. Unlike lotteries or bandits, the quality of one's decision depends upon the decisions made by the other party. Behavioral game theory instead studies decisions in terms of optima and equilibria (Osborne et al., 2004), where a quality move is one with high expected value given the possible choices of the opponent. However, these studies are rarely focused on the risk-taking aspect of decision making, preferring instead to emphasize move quality with respect to an equilibrium or the expected value. Finally, while behavioral games can be immensely useful in generating customized models of local settings, the simplifications they impose on decision making may not generalize to more complex, highly strategic, or highly incentivized domains (Camerer, 1991).

One promising domain in which we can study adversarial risk taking is chess. Chess is a popular lens through which research can study a variety of cognitive processes because the well-defined environment and rules provide a tractable, yet sufficiently complex model to study perception, memory, problem solving and risk taking. Furthermore, chess uses the ELO rating system which is a precise, zero-sum measure of player skill ranging from novices (ELO < 1200) to experts (ELO > 2200). Additionally, algorithms such as Stockfish (*Stockfish: Strong open source chess engine*, n.d.) allow for objective evaluations of move quality and positional advantage that would be difficult or impossible to define in many other naturalistic settings. Finally, chess is a popular game, with game data widely available at massive scales not attainable in a laboratory setting, which allows for analyses across cultures(Chassy & Gobet, 2015), genders (Gerdes & Gränsmark, 2010), ages (Charness, 1981), and skill levels (Chassy & Gobet, 2020). The combination of well-specified evaluation algorithms with difficult, highly strategic game play makes chess an attractive middle ground between naturalistic, adversarial situations and controlled laboratory experiments.

A measure of risk in adversarial games needs to take into account how opponents respond to a particular action, rather than an *a priori* evaluation of that action. Indeed, classic measures of risk in chess rely on the distribution of outcomes observed from professional chess games. In particular, risk in chess has previously been measured based only on the first few moves: whether the opening move choice (Chassy & Gobet, 2015) or the first few moves yielded a relatively high rate of draws (Dreber, Gerdes, Gränsmark, & Little, 2013; Gerdes & Gränsmark, 2010) at the end of the game. The logic of this measure is that a risky opening is one that reduces the probability of draws while increasing the probability of wins as well as losses. This measure may be useful for evaluating the strategy of a professional chess player who picks among

a number of well-studied openings, however, it is predicated on a risk assessment arising from a forecast of a full game trajectory. This measure falls short of our goals in two ways. First, we want to apply a measure of risk to multiple individual moves in a game, so as to enable the analysis of risk seeking behavior as it varies with changes in game state. Second, we want a measure of risk that would represent an organic online calculation — such as may be undertaken by amateur players — rather than a strategy chosen based on *a priori* study of eventual game outcomes by an expert.

To this end, we propose a new algorithm to measure risk in chess that can be evaluated at each individual board-state, providing a measure of risk for each move not available through current measures. Specifically we measure risk of a given move as the variance of expected outcomes after the opponents immediate response. This 1-step lookahead variance is consistent with measures of risk from economics and finance that define risk as the variance around an expected value (Rothschild & Stiglitz, 1970; Damodaran, 2007). The 1-step lookahead measure is not only better aligned with how amateurs evaluate the outcomes of their moves, but also naturally varies from move to move. We applied this measure of risk to over 1.3 billion online chess games across all skill levels, ranging from novices (*ELO* < 1200) to masters (*ELO* > 2200). We present evidence that: (i) risk-taking strategies depend on the skill of the player and change depending on the skill of their opponent, (ii) players are more risk averse when games are close, and (iii) players are more risk-seeking following a blunder.

Methods

Measure of Riskiness

To measure the relationship between current decisions and future outcomes present in the real world, our measure of risk is defined in terms of the opponent's probable responses to a given decision. Specifically, to measure risk of any given move, each possible subsequent board state needs to be evaluated for quality, i.e., whether the resulting position is advantageous for the player, as well as tabulated for frequency, i.e., how often that response is played. With the probability or weight (w_i) and quality (x_i) of possible subsequent game states, we can evaluate the riskiness of a move in terms of the weighted variance of possible outcomes,

$$\sigma^2 = \frac{V_1}{V_1^2 - V_2} \sum_{i=1}^{N} w_i(x_i - \mu^*)^2,$$

where V_1 is the sum of the weights (w_i), V_2 is the sum of the squared weights, and μ^* is the weighted average of game state qualities (x_i). This formulation of weighted variance corrects for the uncertainty arising from potentially undersampling groups of responses (i.e. from not having certain possible moves included in the dataset) (*Statistics*, n.d.), however given the scale of the data, using a frequency weighted variance yields very similar results. Intuitively, risky moves are those that result in a high variance among possible opponent
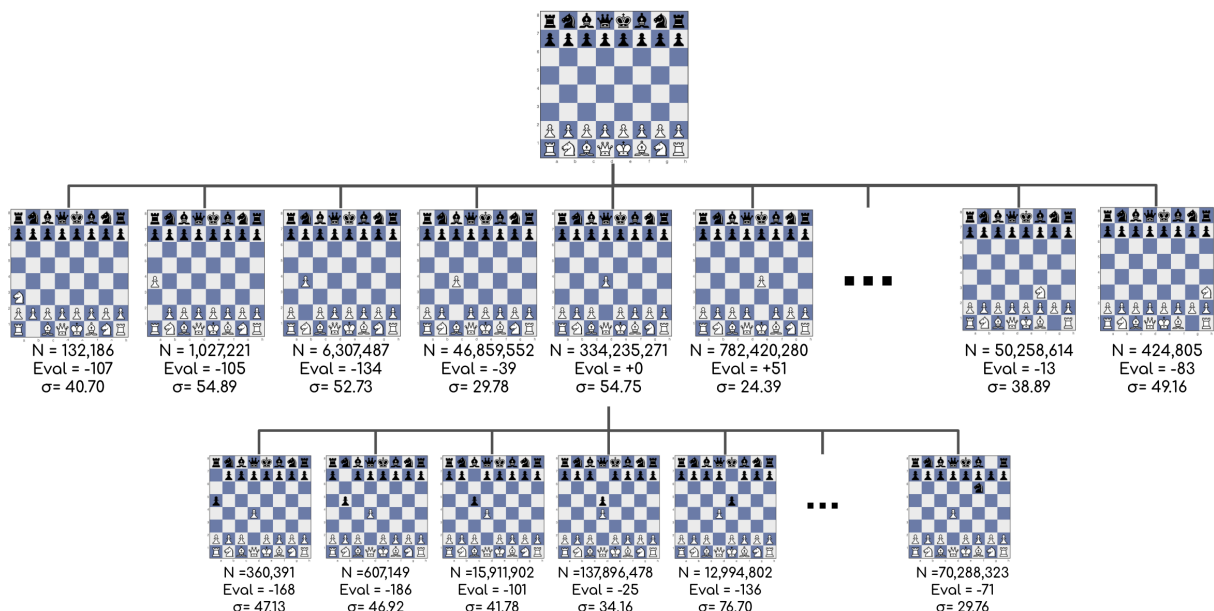
Figure 1: A representation of the move graph. The root node represents the starting board position, with each subsequent board state represented with its own node. Each node stores: the empirical count of times it was played in the dataset, an evaluation of the board state using the Stockfish algorithm, and σ, the standard deviation of possible response moves, weighted by their empirical frequencies.

responses; i.e. moves that could either greatly strengthen or weaken the player's position. Inversely, moves with low variance amongst outcomes, are correspondingly less risky. This method not only leverages classical definitions of riskiness as a measure of variance, but also utilizes the rich empirical distributions of actual past games.

To calculate a weighted variance of possible outcomes, we need to accurately measure the value of each subsequent game state for the player as well as the probability of the opponent choosing that state. To evaluate the quality of a given board state, we rely on the Stockfish evaluation engine which uses advanced alpha-beta pruning search to evaluate a particular board state (*Stockfish: Strong open source chess engine*, n.d.). Since play based on these board-state evaluations allows Stockfish to achieve super-human chess performance, we take them to be the "ground truth" evaluation of the quality of a position. These board state evaluations allow for a measure of game state quality relative to the player (i.e. are they in an advantageous or trailing position) for each state represented in our game-state graph.

To get the probability of each of the possible game states following a particular move, we compute the empirical distribution across all 1-step future states from historical games. However, to be effective, this method requires a volume of data not previously used in other chess research. Chess has an estimated branching factor of around 35, which means in any given game state a player has approximately 35 legal moves. Each move in turn leads to myriad possible countermoves, and the number of possible game-states blows up exponentially, requiring increasingly large amounts of data to have meaningful empirical counts for each board state only

a few moves into the game. To address this, we used the lichess.com online chess database and big data parsing techniques to extract all games played on the site between January 2013 and June 2020. The resulting dataset is several orders of magnitude larger than other recent large-scale studies of human decision making via chess, and dwarfs the amount of data accessible in a traditional lab setting.
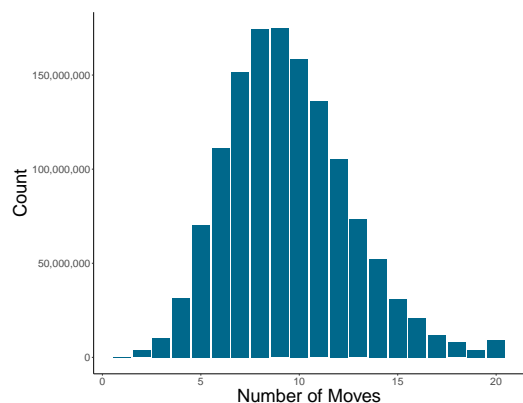

Figure 2: The number of games in the dataset as a function of the game depth (number of moves) represented in the graph. This does not reflect the complete game depth because we retain only game states that had occurred in at least 500 out of the 1.3 billion total games.

Data Collection

We collected 1,337,825,050 classical and Blitz chess games from the open source lichess.org database (*lichess.org open database*, n.d.). These games were read in as .pgn files and parsed to extract pertinent information, such as player ELO

rating, timing data, and the list of moves played. The parsed move sequences were then parsed into a directed acyclic graph with nodes representing unique board states. A given board state was included in the graph if it had been observed in at least 500 games. This constraint yielded a total of 904,563 nodes, or possible game states. Each node consisted of a unique board state, its empirical frequency (number of games in which it appeared), and Stockfish evaluation. We then traverse the graph and calculate the risk of a given move in terms of the variance of game-state evaluations of its immediate "children" (subsequent game states).

Results

With our move graph built and risk calculated for each board state, we can begin to answer questions about how player characteristics and game dynamics affect risk preference. Of particular interest is how decision making changes in this adversarial setting, where each decision (move) is made in consideration of opponent responses, and the current world state. We focus on three key questions 1) Are there expertise effects in risk preference? 2) Are there state-dependent effects in risk-taking? and 3) Do people take more risk following a mistake?

Experts take less risk

Previous research has suggested that experts and novices perceive risk differently, suggesting that while experts have more accurate perceptions of risk (Bostrom, 1997), they may be prone to be more risk seeking in skill-based tasks (Thomson et al., 2004). These results however have not been shown in adversarial games. In fact, a recent study of risk-taking in chess showed the opposite effect; that players with higher ELO ratings played less risky opening moves (i.e., moves that had higher draw percentages), than amateurs (Chassy & Gobet, 2020). This analysis however was limited only to the opening move, so we are interested in whether this effect holds over the course of the game, where the complex dynamics could conceivably influence this result. We grouped players into ELO bins ranging from novices (ELO <1100) to experts (ELO >2100). Figure 3 shows the mean risk taken (log standard deviation) by player ELO. Since ELO is a rating of player skill we see a clear decrease in risk taking by stronger players $[F(9, 984600) = 704.3, p < 0.001]$.

These expertise effects led us to investigate whether players exhibit a consistent risk-taking strategy in all games, or whether they adapted to the strength of their opponent.

Opponent-dependent risk-preferences

A key difference between lotteries and adversarial games is that the source of the risk is determined by the choices of the other party and thus riskiness is opponent-dependent. To explore whether players indeed exhibit differing risk preferences depending on their opponent, we binned players by their ELO and calculated the mean riskiness for each player/opponent ELO combination. Figure 4 shows two main effects: first, as reported above, stronger players are more
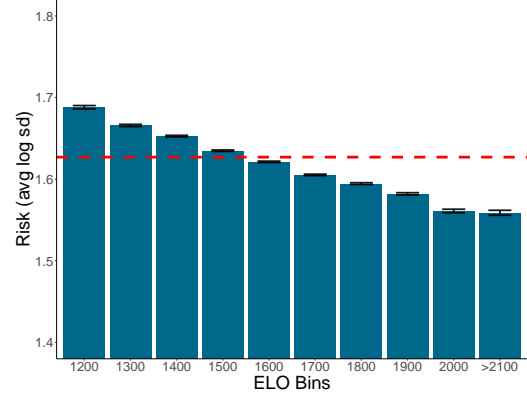

Figure 3: Mean log risk taken by ELO group, with standard error bars shown in black. The average risk across all groups is shown in red.

risk-averse; and second, as the quality of their opponent increases, players, regardless of strength, become increasingly risk-averse. For players with the highest ELO ratings, their behavior seems constant regardless of opponent's rating.

These results suggest that players do change their risk-taking strategy depending on who their opponent is. One possible interpretation is that risk taking is an effort to fool or outmaneuver an opponent (e.g., gambits) and that as the opponent becomes stronger, these outmaneuvering attempts decrease. One caveat is that there are likely to be complex game dynamics that govern the risk decisions of players, especially in games among weaker players who may exhibit "risky" behavior out of erratic incompetence rather than an informed risk-seeking strategy. To further understand how game states may influence risk preference we next explored how the evaluation of the game state (whether a player is winning or losing) influences their tolerance for risky play.
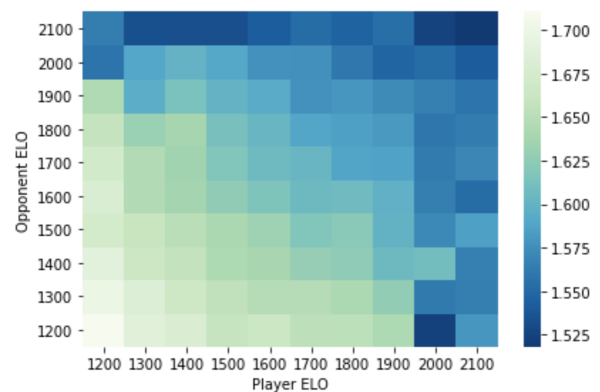

Figure 4: A heatmap of the mean log risk taken by players binned by the player ELO and the ELO of their opponent.

State-dependent risk-preferences

How does risk-taking behavior change as a result of past performance? These effects are commonly studied in reference to gambling behavior, and a variety of effects have been observed that make potentially contradictory predictions (Ayton & Fischer, 2004). For example, the "hot-hand" fallacy

predicts that riskiness increases after a series of successes (Gilovich, Vallone, & Tversky, 1985), and the "house money" hypothesis similarly predicts increased risk when well in the lead (Thaler & Johnson, 1990). These can be contrasted with the gambler's fallacy which predicts that gamblers will be less willing to bet after a big win and more willing to bet following a string of losses since they are "due" to win (Clotfelter & Cook, 1993). Finally the "break-even" hypothesis predicts that players will become more risk averse as they get closer to a neutral winning point (Tversky & Kahneman, 1974; Thaler & Johnson, 1990). A study of professional online poker players explored these competing hypotheses and showed that players play consistent with a "break-even" strategy; taking more risk following big losses and avoiding risk when they approach their break-even point (Smith, Levere, & Kurtzman, 2009).

To determine whether this effect holds in chess, we calculated the value of the game state prior to each move. A positive game score suggests the player is in an advantageous position, while a negative score suggests they are at a disadvantage. Figure 5 shows the mean log risk taken by game score. The average risk taken by players is shown in blue and the maximum risk they *could* have undertaken in their position is shown in grey. While the maximum risk is fairly stable across game scores, the U shaped blue area shows increased risk-taking in games where the player is materially leading or trailing, but relative risk aversion in close games. Essentially the *proportion* of risk undertaken increases in winning and losing positions despite similar maximum possible risks in each position. These results are consistent with the "break-even" findings in poker, with increased risk following losses and risk aversion when scores are close. Interestingly, we also find support for the "house money" hypothesis in which risk-taking increases in winning positions.
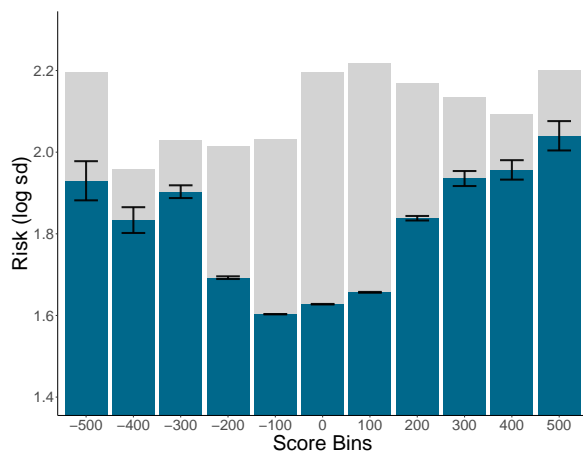

Figure 5: Mean log risk taken by players binned by current game score is shown in blue with black standard error bars. The maximum possible risk (as measured by greatest risk taken in that position in the dataset) is shown in grey. Positive scores indicate the player is in a winning position before making the move.

Increased risk-taking following mistakes

We have shown that risk-taking behavior changes based on the result of previous actions and the current world state, but the previous results do not account for *how* one got into the winning or losing position. One unique aspect of adversarial games is that the game state is determined by the cunning and mistakes of both the player and their opponent. Here we explore whether players increase risk-taking following a big mistake in an effort to compensate for lost ground.

The Stockfish evaluation of a given game state assumes that the opponent will make the best possible move on the next turn. Consequently, if the Stockfish score is 50 for white (-50 for black) after white's move, the best move black could make would yield exactly the same evaluation (-50 for black) after black's move. Any other move for black would yield a worse score. Based on this logic, we can objectively determine the quality of any given move relative to Stockfish, and define a blunder as a move that yields a value at least 100 centipawns below the optimal evaluation.

Figure 6A shows the histogram of moves binned by their difference from the optimal Stockfish move, with the blunders shown in red. Although our binning dedicates most bins to blunders, as a whole they are not especially common ($\approx 15.5\%$). Nonetheless, they represent a large swing in the quality of one's position. Figure 6B shows that following a move in which the player commits a blunder, the subsequent move is on average riskier than moves played following non-blunders [$F(9, 929013) = 1627, p < 0.001$], despite risky moves being *available* across all bins (shown in grey). To address the fact that weaker players are more likely to commit blunders, and blunders are more likely to happen later in the game, we conducted multiple regression to control for these factors. We found that even when controlling for move number and player ELO, moves are significantly more risky following blunders ($\beta = 0.101, p < 0.001$).

Discussion

The ability, or perceived ability, to control how adversaries respond in games make the study of risk taking in these settings fundamentally different than games with a purely mechanical source of risk, such as lotteries. The current work uses a large-scale dataset of chess games, and novel measure of risk taking, to explore how past performance, expertise, and the current world state influence risk taking in strategic, adversarial games. Specifically, we show that players take more risk when in clearly winning or losing positions, and prefer a more risk-averse strategy when games are close. These findings are consistent with "house-money" and "break-even" accounts of risk-taking behavior (Smith et al., 2009). We also show that expert players prefer a more risk-averse strategy compared to novices; an expert–novice difference that may not be present in games with mechanistic risk.

In addition to state-dependent effects, we also find opponent-dependent effects on risk-taking behavior. Consistent with (Chassy & Gobet, 2020), we find that players' pref-
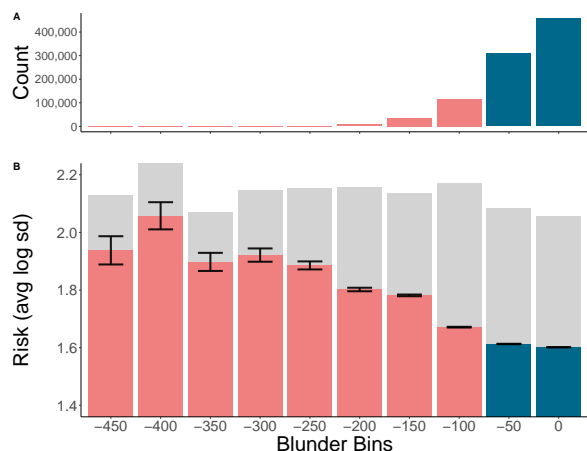
Figure 6: Below: Mean log risk taken binned by quality of previous move. A blunder (in red) denotes a move that was evaluated to be at least 100 centipawns below the optimal Stockfish evaluation, non-blunders are in blue. The maximum possible risk (as measured by greatest risk taken in that position in dataset) is shown in grey. Above: histogram of frequency of each bin; blunders occur approximately 15.5% of the time.

erence for risk depends on who they are playing. In future analyses we plan to explore the extent to which these risk taking patterns represent meaningful strategies versus uninformed poor play. In particular we are interested in studying in what circumstances risky play is *optimal*, reflecting an strategic weighing of risk and reward, and the extent to which it correlates with win percentage. While these results are preliminary and will require further investigation into the complexities of chess games and why these results arise, they offer a promising future direction for cognitive scientists trying to understand how people reason about risk in different settings.

Chess provides a valuable lens through which we can study risk taking in adversarial settings. The combination of objective world state evaluations, reliable player competency measures, and massive amounts of available data make it a tractable, yet complex research vehicle. We leveraged these features to develop a novel measure of risk-taking that is in line with definitions in economics and behavioral finance. This new method for calculating risk in chess, combined with the wide accessibility of chess data from a variety of sources means this work could be extended to study a variety of otherwise intractable questions in areas beyond chess, such as cultural or gender differences in risk taking. Future work should also explore the extent to which other key findings in stochastic risk settings are mirrored in adversarial settings, and the extent to which these results generalize to other adversarial settings such as negotiations.

Acknowledgements

References

Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & cognition*, *32*(8), 1369–1378.

Bostrom, A. (1997). Risk perceptions: experts vs. lay people. *Duke Envtl. L. & Pol'y F.*, *8*, 101.

Brockbank, E., & Vul, E. (2020). Recursive adversarial reasoning in the rock, paper, scissors game.

Camerer, C. F. (1991). Does strategy research need game theory? *Strategic Management Journal*, *12*(S2), 137–152.

Charness, N. (1981). Search in chess: Age and skill differences. *Journal of Experimental Psychology: Human Perception and Performance*, *7*(2), 467.

Chassy, P., & Gobet, F. (2015). Risk taking in adversarial situations: Civilization differences in chess experts. *Cognition*, *141*, 36–40.

Chassy, P., & Gobet, F. (2020). Influence of experts' domain-specific knowledge on risk taking in adversarial situations. *Journal of Expertise/June*, *3*(2).

Clotfelter, C. T., & Cook, P. J. (1993). The "gambler's fallacy" in lottery play. *Management Science*, *39*(12), 1521–1525.

Damodaran, A. (2007). *Strategic risk taking: a framework for risk management*. Pearson Prentice Hall.

Dreber, A., Gerdes, C., Gränsmark, P., & Little, A. (2013). Facial masculinity predicts risk and time preferences in expert chess players. *Applied Economics Letters*, *20*(16), 1477–1480.

Gerdes, C., & Gränsmark, P. (2010). Strategic behavior across gender: A comparison of female and male expert chess players. *Labour Economics*, *17*(5), 766–775.

Gilovich, T., Vallone, R., & Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive psychology*, *17*(3), 295–314.

Harrison, G. W., Rutström, E. E., et al. (2008). Risk aversion in the laboratory. *Research in experimental economics*, *12*(8), 41–196.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in cognitive sciences*, *13*(12), 517–523.

Kahneman, D., & Tversky, A. (2013). Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part i* (pp. 99–127). World Scientific.

lichess.org open database. (n.d.). `https://database.lichess.org/`. (Accessed: 2020-06-30)

Mata, R., Frey, R., Richter, D., Schupp, J., & Hertwig, R.

(2018). Risk preference: A view from psychology. *Journal of Economic Perspectives*, *32*(2), 155–72.

Meyer, R. J., & Shi, Y. (1995). Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem. *Management Science*, *41*(5), 817–834.

Osborne, M. J., et al. (2004). *An introduction to game theory* (Vol. 3) (No. 3). Oxford university press New York.

Rothschild, M., & Stiglitz, J. E. (1970). Increasing risk: I. a definition. *Journal of Economic theory*, *2*(3), 225–243.

Slovic, P. E. (2000). *The perception of risk.* Earthscan publications.

Smith, G., Levere, M., & Kurtzman, R. (2009). Poker player behavior after big wins and big losses. *Management Science*, *55*(9), 1547–1555.

Statistics. (n.d.). Retrieved from `https://www.gnu.org/software/gsl/doc/html/statistics.html#mean-standard-deviation-and-variance`

Stockfish: Strong open source chess engine. (n.d.). `https://stockfishchess.org/`. (Accessed: 2020-06-30)

Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management science*, *36*(6), 643–660.

Thomson, M. E., Önkal, D., Avcioğlu, A., & Goodwin, P. (2004). Aviation risk perception: A comparison between experts and novices. *Risk Analysis: An International Journal*, *24*(6), 1585–1595.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, *185*(4157), 1124–1131.

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*, *106*(4), 1039–1061.