

Gerdes, Christer; Gränsmark, Patrik; Rosholm, Michael

Working Paper

Chicken or checkin'? Rational learning in repeated chess games

IZA Discussion Papers, No. 5862

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Gerdes, Christer; Gränsmark, Patrik; Rosholm, Michael (2011) : Chicken or checkin'? Rational learning in repeated chess games, IZA Discussion Papers, No. 5862, Institute for the Study of Labor (IZA), Bonn,
<https://nbn-resolving.de/urn:nbn:de:101:1-201107283903>

This Version is available at:

<https://hdl.handle.net/10419/51900>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 5862

**Chicken or Checkin’?
Rational Learning in Repeated Chess Games**

Christer Gerdes
Patrik Gränsmark
Michael Rosholm

July 2011

Chicken or Checkin'?

Rational Learning in Repeated Chess Games

Christer Gerdes

*SOFI, Stockholm University
and IZA*

Patrik Gränsmark

SOFI, Stockholm University

Michael Rosholm

*Aarhus University, CIM
and IZA*

Discussion Paper No. 5862

July 2011

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Chicken or Checkin'? **Rational Learning in Repeated Chess Games***

We examine rational learning among expert chess players and how they update their beliefs in repeated games with the same opponent. We present a model that explains how equilibrium play is affected when players change their choice of strategy when receiving additional information from each encounter. We employ a large international panel dataset with controls for risk preferences and playing skills whereby the latter accounts for ability. Although expert chess players are intelligent, productive and equipped with adequate data and specialized computer programs, we find large learning effects. Moreover, as predicted by the model, risk-averse players learn substantially faster.

JEL Classification: C73, D83

Keywords: rational learning, risk aversion, beliefs

Corresponding author:

Patrik Gränsmark
Swedish Institute for Social Research
Stockholm University
SE-10691 Stockholm
Sweden
E-mail: patrik.gransmark@sofi.su.se

* We are grateful for comments by Åsa Rosén, Anders Stenberg, Eskil Wadensjö, Robert Östling, and the participants at the SOFI seminar. Moreover, we have benefited from comments on chess-specific issues from chess Grandmaster and reigning Danish champion Allan Stig Rasmussen and chess FIDE master and former national Danish coach Per Andreassen. Thanks to Susan Stilling, who did a very careful proof-reading of the final document.

1 Introduction

Imagine two kings at the opposite sides of a battle field, each one in the lead of an army of knights and archers. Each king is considering his options: should he attack with the risk of losing or should he offer a truce? He knows that if he offers a truce, he signals potential weakness and that he would prefer not to engage in a battle, hence, offering a truce could encourage the opponent to attack. On the other hand, perhaps both kings prefer a truce but if none of them dares to offer it, they will have to take their chances on the battle field. However, had the two kings agreed to a truce in the past, they would probably feel a lot less uneasy. They would attach a higher probability to the possibility of establishing a new truce, which in turn increases the probability that one of the kings will propose it. When such a peaceful outcome has occurred several times in the past, the kings will probably reduce the size of the army; why keep an expensive reserve when a war is not very likely? In this example, the kings have learnt from past encounters, and this new information affects the future choice of strategy. Parallels can be drawn to the financial market: When there has been a certain amount of transactions between two traders, lenders and borrowers learn that the risk of default is not very high, so they are willing to reduce their risk premium. The problems for the kings, lenders and borrowers is that if there is some kind of (supply) shock, they will be more vulnerable, since they have lower reserves than before the learning process.¹

Although these learning processes are important and common, little is known about how they work in practice. This is particularly true for *rational* learning. To address this issue, we turn to real chess games, performed by expert chess players. The purpose of the present paper is to examine strategic learning in repeated Bayesian games, using a large real world field data set. We test whether expert chess players, in repeated chess games with the same opponent, obtain additional information which leads to an update of their beliefs regarding the opponent's degree of risk aversion. In addition, we show that risk-averse players learn and update differently than risk-loving players. Specifically, we test whether risk-averse players, when they meet repeatedly, become increasingly likely to end the game in a so-called 'arranged draw', which is an outcome the players can agree upon at any time during the early phases of the game while it is still undecided on the board, and where they

¹ To give some further examples from economics, consider a union that must decide whether to accept the wage offer from the employer or to begin a strike with a risk of losing even more. Should a market-leading company let a newcomer enter the market and share the profit or should it start a price war with the risk of having to let the competitor in anyway?

share the points. Thanks to the tremendous development in the accessibility of chess data over the last decade, we exploit the existence of a large panel data set, where the players have an almost perfect record of the game history of their opponents, including an evolving measure of playing skills of all players.

It is known from experimental behavioral economics that strategic learning occurs, see Camerer (2003) and Young (2004). Recent theories model rational learning in repeated Bayesian games, but empirical testing of these models still lags behind, however, and existing studies are largely based on small-scale lab data. One reason for this is that it is difficult to construct beliefs that reflect genuine uncertainty about the opponent's strategies, yet at the same time, narrow enough to permit learning and to be able to observe the learning process in practice, i.e. it should not require too many repetitions for learning effects to emerge. To find field data that comply with these conditions has proven difficult.

Most empirical research in behavioral game theory is based on lab experiments which have been criticized for being too unrealistic, producing results that are valid only for particular subgroups in a sterilized setting. The number of observations is usually rather low. On the other hand, many lab experimentalists are not satisfied with findings resulting from field or observational data as it is typically not possible to control for confounders satisfactorily. Theorists often desire a continuous line of vertical logic which is almost impossible to achieve in large-scale empirical studies. In short, there is a trade-off between having control over the experiment on the one hand and the number of observations and population representativity on the other.² In this paper we bridge the gap between small-scale lab experiments and large-scale imperfect data by employing large-scale field data collected from real chess games but in a very controlled setting, where the rules of the game are the same all over the world, and where it is possible to control for exact playing skill by exploiting the Elo (1978) rating system (see section 2.2). The rules and homogeneity of the game of chess offer a setting that is a step towards a controlled lab experiment but at the same time supplies a data set reflecting real behavior with 1.5 million observations in a panel data structure.

We contribute to the literature by bringing chess data into the analytical toolbox of empirically oriented economists and by showing that rational learning in games occurs in large-scale field data. In addition, we present a model that explains how players, in a

² Levitt and List (2007) discuss the advantages and drawbacks of lab experiments. They argue that having one's actions scrutinized by others may have an unintended influence on the lab participants. For a discussion in favor of lab data, see Falk and Heckman (2009).

Bayesian game, change their choice of strategy as they receive additional information about the opponent after each encounter. The model also predicts that risk-averse players learn faster, which is supported by our empirical findings. Furthermore, we show that the likelihood of a risk-averse outcome increases continuously, when two risk-averse players meet repeatedly, although they knew the general risk type of the opponent in advance. Our findings suggest that equilibria may evolve over time as the learning process proceeds. The mere duration of stability may affect the level of risk taken by individuals in the sense that there is little need for a margin if the risk of default is very small.

The fact that elite chess players are considered to be among the most intelligent subpopulations, with high levels of cognitive ability, has attracted several researchers within the field of economics. Palacios-Huerta and Volij (2009) and Levitt et al. (2009) study chess grandmasters to analyze how they use backward induction to solve the centipede game. The former study finds that grandmasters do indeed use perfect backward induction, whereas the latter finds little support for that. Moul and Nye (2009) find that players from the former Soviet Union could improve their result by agreeing to early or pre-arranged draws. Simon (1955) contributed to economics with an influential paper, where he modeled the rational choice of chess players. These studies have in common that they take advantage of the fact that elite chess players constitute an upper boundary of the population, as far as rationality is concerned, and thereby serve researchers to establish a benchmark.³ We have chosen to study the behavior of expert chess players for the same reason and due to the fact that the game setting is orderly and the information set, both for the players and the researchers, is very rich, which reduces the potential impact of confounders.

The paper is organized as follows. The next section provides a chess background, and section 3 discusses risk-taking and the measurement of risk aversion in chess. Section 4 discusses the conceptual framework of learning in repeated games and presents the theoretical model. In section 5 we discuss the data and the econometric strategy. Section 6 contains the results of the estimations, the sensitivity analyses and a placebo analysis. Section 7 concludes.

³ For other references to studies on chess players, see, for instance, Gobet (2005), Roring (2008) and Ross (2006).

2 A Brief Introduction to Chess

This section contains a brief introduction to chess, a discussion of the opening strategies, an intuitive description of the Elo (1978) rating system (a formal description is given in the Appendix), and ends by discussing the information set of a chess player when preparing for a game.

2.1 The Chess Game

Chess is a sequential game where the players make moves in turn with white and black pieces on the chess board with the aim of capturing the opponent's king. There are three possible outcomes of a chess game: you win, you draw (a tie), or you lose, which scores 1 point, $\frac{1}{2}$ point, and 0 points, respectively.

The fact that there is a third outcome, a draw, and the fact that it can be agreed upon by the players at any time during the game, makes chess suitable for studying risk preferences.⁴ The offer is valid and binding until the opponent has made his next move. There are no rules regulating the minimum number of moves that have to be played before the players can agree upon a draw, except that the game must have started. It is considered impolite to offer a draw more than once during a game.

Most chess games recorded in Chessbase (which is described in sub-section 2.4) are played in chess tournaments or in team tournaments. In chess tournaments, it is normal to play one game per day, and a tournament may consist of five to thirteen rounds, with ten being the norm. Since each player has an allotted time that is not to be exceeded, a game lasts at most six to seven hours.⁵ If your time limit is exceeded, you lose the game (if the opponent still has material enough to capture your king, otherwise it is a draw).

2.2 The Opening

The description of the opening strategy is important for our definition of risk-taking in the next section, hence, we devote some time to explaining the opening strategy.

There are different ways in which a player can affect the course of the game. The most obvious way is through the choice of opening. At the beginning of a chess game, both players choose an opening strategy (a strategic development scheme for their pieces) that will steer the game towards a style of play that best suits them and at the same time makes life less

⁴ Only the player, whose turn it is to move, may offer a draw. He has to make his move and then make the offer of a draw immediately after the move has been made.

⁵ A chess clock has two clocks, one for each player.

comfortable for the opponent. All expert chess players have a prepared set of opening strategies to be used in different situations. One's chosen set of openings is called an "opening repertoire" (OR). To optimize performance, a considerable effort is dedicated to creating an opening repertoire that matches one's personality. In the opening, the player must make a decision about the task of each piece, on which square to put it, whether its main purpose is to attack or defend, etc.

When the game starts, it is in a balanced position, signifying that there is no advantage for either side. The opening moves are *theoretical* in the sense that they have been worked out beforehand as to maximize the expected score for each side. The theoretical moves are to a large extent memorized before the game, and, since no calculation is needed, these moves are often played quickly. While the game is still in the theoretical phase, the position remains balanced. If the theoretical extent of a certain opening variation lasts for twenty moves, then the "real" play and a potential deviation from the saddle point equilibrium does not start until the 21st move.

Chess opening strategies have been analyzed in extreme detail both by chess players and by computer programs. All expert chess players nowadays use computer programs and chess databases. The chess database reports the relevant statistics for each opening strategy, for instance, the mean score for each opening (based on the games stored in the database). This means that a player easily can observe how well an opening has scored in the past and the distribution of outcomes. Computers have the advantage that they can calculate much faster and more accurately than the human brain. For every calculated variation, the computer program evaluates the position and expresses it in units of pawns (e.g. +1 usually means that the position corresponds to white being one pawn up, materially). When both players play optimally, this value is (close to) zero in equilibrium.

Since there are quite a large number of possible strategies by the opponent for each additional move, the opening theory is limited (by cognitive constraints). Most opening strategies cover reasonable moves made by the opponent for approximately the first 15 to 20 moves.⁶

2.3 The Elo Rating System

A landmark for establishing chess as an analytical tool was the introduction of the Elo rating system (Elo, 1978), which made it possible to compare the strength of chess players on a

⁶ See de Firmian (2009).

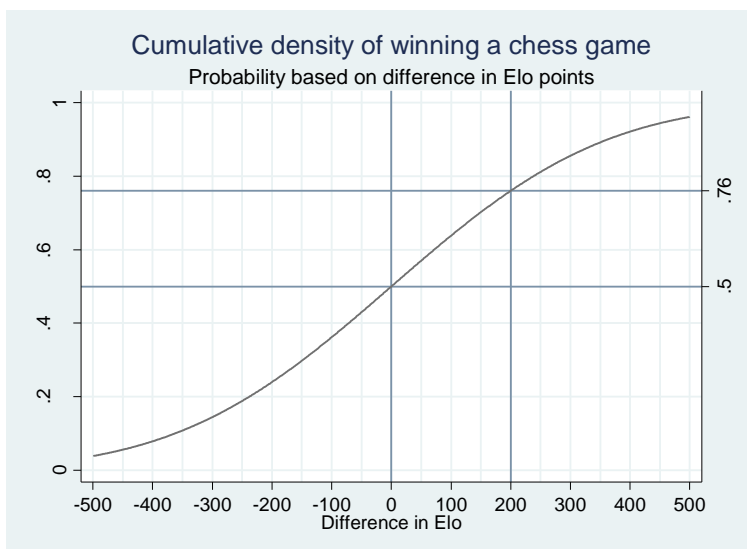
metric scale.⁷ Named after its inventor, Arpad Elo, it has become the benchmark rating in chess.

“[It] provides chess researchers with a valid measurement device unrivalled in other areas of expertise research. It is a true gold standard in individual-difference research.” (Charness 1992, p. 6).

Thus, with reference to Elo (1978), it has become possible to measure skills on objective grounds, there are no “subjective assessments” (Chabris and Glickman 2006, p. 1040).

The Elo rating scale is constructed in such a way that the Elo difference between two players corresponds to an exact expected score (performance). The Elo rating of a player increases, when the player scores above the expected score and vice versa. Figure 1 displays the scoring probabilities for different relative Elo ratings. In the Appendix we show how the Elo rating and scoring probabilities are computed.

Figure 1 Expected scores for varying Elo differences.



The expected score between two equally skilled players is .5, i.e. when the Elo difference is zero the expected score is 50 %. If a player has an Elo rating of 200 points more than the opponent, the expected score is about .75, see Figure 1.

⁷ The history of the Elo rating system is described in Ross (2007).

2.4 The Information Set of a Chess Player

A chess player in a tournament or team tournament typically knows in advance his next opponent, including the Elo rating, and will spend a considerable amount of time (from a few hours to half a day) preparing for that specific game, studying the opponent's style of play, opening strategies, etc. In doing so, the player uses books, computer programs, and a database of chess games. Chessbase is arguably the most comprehensive and most used of such databases, and it contains a total of more than five million chess games. For each game, it contains all the moves by the players, the date of the game, the names of the players, their Elo rating at the time of the game (or at any point in the past), and the outcome. Hence, when confronting a new opponent, a chess player has the possibility to study the playing skill (the Elo rating), the opening strategies and playing styles of the opponent, his strengths and weaknesses as reflected in his past games. In addition, the players have information about the age, nationality and gender of the opponent. The Chessbase database also calculates the score of different opening strategies based on real games, and the expected score based on computer evaluations of the positions.

In short, the chess player, when preparing for the next game, has – except for potential private information via e.g. friendship - access to exactly the same information set a priori as the econometrician who can access the same database.

3 Risk-taking in Chess

One standard way of defining risk aversion in economics is that a risk-averse agent prefers the utility of the expected value rather than the expected utility. Let p , q and r denote the probability of a win, draw, and loss, respectively. Then, risk aversion in chess can be formally stated as

$$U(p \cdot 1 + q \cdot .5 + r \cdot 0) > p \cdot U(1) + q \cdot U(.5) + r \cdot U(0) \quad (1)$$

For a risk-loving agent the inequality is reversed.

The relation between p and r depends on the Elo difference between two given players. More precisely,

$$p = r + [\varphi - (1 - \varphi)] \quad (2)$$

where $\varphi \in (0, 1)$ is the expected score of the player in focus, that is, $\varphi = p \cdot 1 + q \cdot .5 + r \cdot 0$. If two players are equally skilled, $\varphi = 1/2$ and $p = r$. However, this still leaves the player with a choice of risk level, since she can choose between riskier openings that provide a large

winning probability (e.g. $p=0.4$, $r=0.4$, $q=0.2$) or a risk-averse opening leading to a smaller winning probability ($p=0.2$, $r=0.2$, $q=0.6$).⁸

If the Elo difference is positive (the player in focus is superior), then $[\varphi - (1 - \varphi)]$ is positive and, consequently, $p > r$, i.e. the probability of a win is greater than that of a loss, and vice versa if negative, but the choice of risk strategy remains. Since the sum of the probabilities must equal one, we can substitute for r in (2) above and obtain $p = \varphi - \frac{1}{2}q$. This is the player's constraint when choosing the opening strategy and the level of risk. Thus, a player's risk preferences affect the choice of p , and all parameters are pinned down when p is chosen. It follows that $p \in [0, \varphi]$, $q \in [0, 2(\varphi - p)]$ and $r \in [p + 1 - 2\varphi, 1 - \varphi]$. To increase the winning probability, the player must accept an increased probability of a loss.

To see that a risk-averse player prefers a draw to playing the game, consider two equally skilled players and then plug in the extreme values of p . Setting $p=.5$ and $q=0$ renders either the score 1 (a win) or 0 (a loss) with the expected score .5. Setting $p=0$ and $q=1$ renders the same expected score $\frac{1}{2}$ with certainty. If both players are risk-loving, they will maximize p ; if both are risk-averse, they will maximize q . If there is one risk-loving and one risk-averse player, then there will usually be an interior allocation of p and q since they are pulling the game in opposite directions in terms of riskiness.

For each additional move that is played in a game, there is a risk of deviation from the theoretical opening balance, with a reduced probability of a draw. Hence, if two equally skilled players have risk-averse preferences, they could reduce the risk by agreeing to a draw at an early stage. If one of two equally skilled players has made a mistake on the board, the probability of a draw is lower than when the position is still in the opening balance. Draws agreed to while still playing theoretical moves are usually referred to as *arranged draws*. Since the players have not really started to play, arranged draws depend on the players' preferences, while draws agreed to at a later stage, when the theory has ended, depend increasingly on the position on the chess board for each additional move.⁹ Most opening

⁸ For a more detailed example, see Appendix.

⁹ Regarding the difference between draws in general and arranged draws, Moul and Nye (2009) write: "Hard-fought games that end in draws are more likely to last longer than collusive or pre-arranged draws. The latter are more likely to be agreed to at an earlier stage when the position on the board is still not fully resolved and it is not clear that one player should win. At a later stage the likelihood is much greater that the position will clearly favor one or the other player." (p.14). Risk preferences affect the preference for an arranged draw and may vary from game to game, depending on the opposition or other circumstances. Drawing preferences may also vary due to temporal variation in health condition, state of alertness/tiredness or if a draw would suffice to obtain a particular objective as for instance the win of a tournament. Naturally, a superior player would have lower drawing preferences against an inferior opponent, while the inferior player would have higher drawing preferences compared to when the players are equal in playing strength.

theory in chess lasts for about 15-20 moves and we have chosen to define an arranged draw as a draw in less than twenty moves.¹⁰

Definition 1. We define an arranged draw (AD) as a game that ended in a draw while the position was still in the theoretical phase. We assume that a game that has ended in a draw in 1-19 moves is an AD. The utility of an arranged draw is denoted $U(\frac{1}{2})$.

Definition 2. Let $U_j(\varphi)$ be player j 's utility of the expected score $\varphi_j \in (0,1)$, where $\varphi = p \cdot 1 + q \cdot .5 + r \cdot 0$, and let EU_j be the expected utility $p \cdot U(1) + q \cdot U(.5) + r \cdot U(0)$, given φ . We define a risk-averse player as having $U_j(\varphi) > EU_j$, a risk-neutral player $U_j(\varphi) = EU_j$ and a risk-loving player $U_j(\varphi) < EU_j$. Furthermore, we say that a player is superior if $\varphi > \frac{1}{2}$ and inferior if $\varphi < \frac{1}{2}$. Finally, a player has a preference for an AD if $U_j(\frac{1}{2}) > EU_j$, i.e., the utility of $\frac{1}{2}$ points is greater than the expected utility, given φ , and we say that this player is of Type I. A player prefers to play the game through if the inequality is reversed, and we say that this player is of Type II.¹¹

The intuition behind Definition 2 is that a type I player prefers a draw given the expected score whereas a type II player prefers to play for a win with the risk of losing given the expected score. Hence, the risk preferences are player-specific while the type I/II categories are game-specific, i.e., the risk preferences depend only on the individual whereas the type depends on the risk preferences and the relative playing skills of the players.

Proposition 1. Assuming that φ_j is distributed such that $E[\varphi_j] = \frac{1}{2}$, i.e. on average a player is as often superior as inferior, then a risk-averse player is more likely to prefer an AD (being of Type I than a risk-loving player, and vice versa.

Proof. The definition of risk aversion is $U(\varphi) > p \cdot U(1) + q \cdot U(.5) + r \cdot U(0)$. Substituting for p and q , we obtain $U(\varphi) > (r + 2\varphi - 1) \cdot U(1) + 2(1 - r - \varphi) \cdot U(\frac{1}{2}) + r \cdot U(0)$ Taking the

¹⁰ Although this definition may seem somewhat arbitrary, it does not affect the results more than marginally. A sensitivity analysis is carried out with the definition 1-15 and 1-23 with similar results.

¹¹ For a player to prefer an AD, the following inequality must hold, $\frac{1}{2} - \frac{r}{2} - \frac{r}{2[U(1) - U(\frac{1}{2})]} \cdot U(0) > \varphi$

expected value of both sides, we have that $EU(\varphi) > E(r + 2\varphi - 1) \cdot U(1) + E2(1 - r - \varphi) \cdot U(\frac{1}{2}) + Er \cdot U(0)$, which is the same as $U[E(\varphi)] > (r + 2E(\varphi) - 1) \cdot U(1) + 2(1 - r - E(\varphi)) \cdot U(\frac{1}{2}) + r \cdot U(0)$. Since $E[\varphi] = \frac{1}{2}$, we know that $U(\frac{1}{2}) > r \cdot U(1) + r \cdot U(0) + (1 - 2r) \cdot U(\frac{1}{2})$ which is equivalent to $U(\frac{1}{2}) > \frac{1}{2}U(1) + \frac{1}{2}U(0)$. The last expression is the very definition of risk aversion, whereas the left-hand side of the inequality is the definition of the utility of an AD. ■

3.1 Measuring Risk-taking: The Openings Classification System

As we have already argued, risk preferences are reflected in the preferences for certain outcomes. The opening strategies are means to increase the probability of a certain outcome occurring.

There exists a standardized classification of 500 chess opening strategies, which are mutually exclusive and exhaustive, called the ECO classification. To create a risk measure, we have categorized each of these 500 openings as either risk-loving, risk-neutral or risk-averse. To obtain such a categorization for each opening (and for each color, white and black), we consulted eight chess experts of different skills with Elo ratings ranging from 2000 to 2600, five men and three women, and asked them to give their characterization of each of the 500 ECO codes.¹² They were instructed to define each opening as *risky*, *neutral* or *safe*. We then compared the opinions of the experts and subsequently we define an opening to be risk-loving, risk-neutral or risk-averse if at least six out of eight experts agreed.¹³ In cases when there were five or fewer votes for either risk-loving or risk-averse, the opening was considered to be risk-neutral. As a result of our experts' assessments, there are two labels for each game, one for each player. We will refer to this risk measure as the *OR risk* preference measure. We will also calculate the ratio of risk-averse (-loving) opening strategies divided by all games played by a player.

In addition, we create a ratio of the number of arranged draws relative to the number of all games played by a player. We refer to this as the *AD risk* preference measure. The OR and AD risk measures are highly correlated as regards risk preferences.

¹² According to the International Chess Federation (FIDE), a player is regarded as an expert if he/she has an Elo rating of 2000 or more. The lowest level required to obtain a Master title is a rating of 2300. A Grandmaster title usually implies an Elo rating of over 2500. In the year 1999 Garry Kasparov reached an Elo rating of 2851, the highest Elo rating ever obtained by a human player.

¹³ The reason for using only 3 categorical values rather than, say, 10, is that the classification requires a very high level of expertise from those doing the categorization. Each of these players needed several hours to complete the survey on the 500 opening codes.

4 Conceptual Framework

4.1 Theoretical Model

There is a first-mover *disadvantage* in chess when proposing an arranged draw. This follows as offering a draw signals potential weakness. The opponent can then choose a strategy that is more efficient. There is also an increased psychological pressure on the player that has revealed potential weakness. For these reasons, players do not want to offer a draw if it is likely that the opponent will reject the offer. Therefore, a player can benefit from learning about the drawing preferences of the opponent before offering an arranged draw. In the present paper, strategic learning is measured by analysing whether the probability of an arranged draw between two players in earlier periods (lagged dependent variables) increases the probability of a future arranged draw. The hypothesis is that when players have agreed to earlier arranged draws, there is less uncertainty and thereby less risk that your opponent will take advantage of you signalling potential weakness. Hence, a risk-averse player has a larger preference for an arranged draw, but at the same time there is a risk involved in trying to avoid risk by offering an arranged draw.

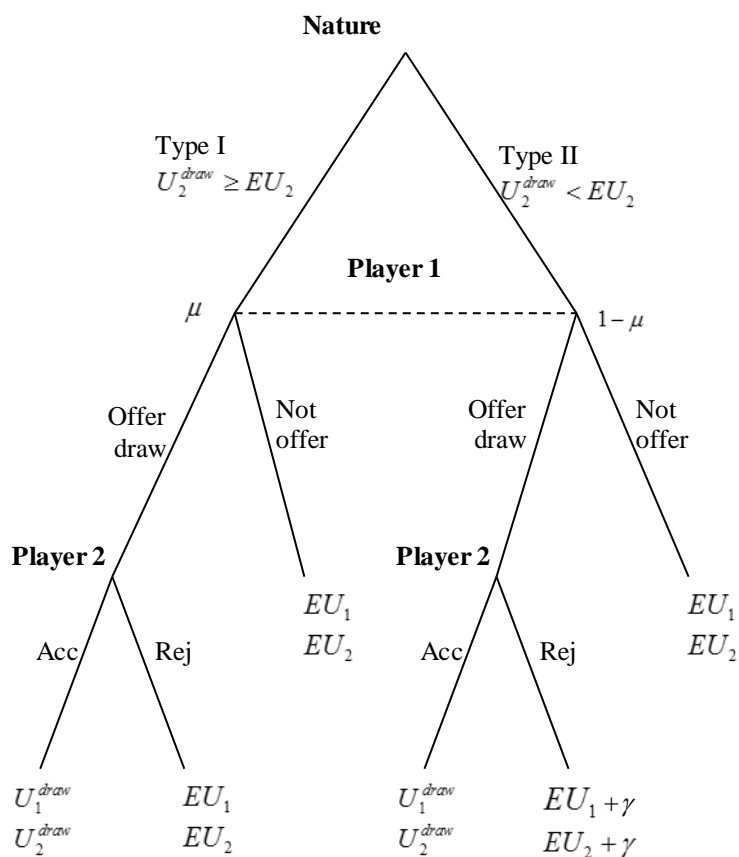
In this game, player 1's strategy set is *offering a draw* or *passing* (not offering a draw) whereas player 2's strategy set is *accepting* or *rejecting* given that an offer has been made. If player 1 passes, player 2 "becomes" player 1 and the game continues. Moreover, when a player has had a draw offer rejected, the game ends (though the chess game continues).

Figure 2 displays the game which starts with *nature* selecting a type for each player, either type I where $U_1(\frac{1}{2}) > EU_1$, or type II where $U_1(\frac{1}{2}) < EU_1$. Since the players have perfect information about the game history of the opponents, they also know the general type of the opponent. If nature picks type II for player 1, then the game stops since a player that prefers to play for a win with the risk of losing rather than taking half a point with certainty will not consider offering a draw as a draw offer is binding. For simplicity, we leave this scenario (nature selecting type II for player 1) out of the game since the choices are then trivial. Note that the game is asymmetric in the sense that a type II player will never act as player 1. As we will see below, this asymmetry affects the optimal play differently for risk-averse and risk-loving players.

When nature has selected type I for player 1, the game in Figure 2 starts. Nature also picks a type for player 2, but the type in this specific game is private information for player 2,

that is, player 1 cannot observe the type of player 2 (symbolized by the dashed line at player 1's information set). Nevertheless, player 1 has a belief about the opponent's utility of a draw. If player 1 chooses to offer a draw, then player 2 can choose either to accept or to reject the offer. If player 2 is of type 1, she will be better off by accepting the offer but should reject if she is of type II. If no offer is made, the payoff is simply the expected utility. If player 1 is of type I and chooses to offer a draw when player 2 is of type II and rejects the offer, a cost γ , where we assume that $\gamma_i \in (0, \varphi_i)$, falls upon player 1 for giving away private information (potential weakness) whereas player 2 receives the reward γ .

Figure 2 – The game of arranged draws in extensive form.



Since it is costly for player 1 to signal that she has preferences for a draw although she does not, player 2 knows with certainty the preferences of player 1 so player 2 simply accepts the offer if she is of type 1 and rejects if she is of type II. Player 1's choice, however, depends on her belief about player 2's type. She will be indifferent between the two strategies when they render the same payoff, that is, when $\mu U^{draw} + (1-\mu)(EU - \gamma) = EU$. This means that player 1

should offer a draw if her belief $\mu \geq \frac{\gamma}{U^{draw} - EU + \gamma}$, and not offer a draw if the inequality is reversed.

The basic idea of the model is that the more information a player receives about the opponent, the better the possibility to determine her utility of a draw. We assume that each player has a certain utility of a draw in a particular game and that a player tries to infer the opponent's utility by interpreting signals. With new information the player updates her beliefs before the next meeting. A player who receives no signals will infer the expected U_j^{draw} of the opponent j to be φ_i (i.e., the stronger is player i , the higher the drawing preferences of opponent j). Each time two players agree to an arranged draw, they receive the signal $\bar{\theta}_i^t$ about opponent i at time t , where $\bar{\theta}_i^t \in \{I, II, \Theta\}$ and I=accepted draw offer, II=rejected draw offer, and Θ =empty set of signals, i.e., no draw offer has been made. Since the game history of all players is common knowledge, the distribution by nature between the two types is also common knowledge. As the number of signals increases, it pulls the posterior beliefs away from the prior mean toward the true value. When the number of signals approaches infinity, the player can infer the U^{draw} of the opponent perfectly.

Let μ_i^1 be player i 's initial belief that the opponent is of type I at time $t=1$. The prior μ_i^1 is a function of the opponent's game history, i.e. $\mu_i^1(\lambda_j^0)$, where $\lambda_j^t \in [0,1]$ is the game history until time t .

The first time two players meet they use the game history of the opponent to infer the type. Thus, their initial belief is $\mu_j^1 = E(\theta_i^1) = \Pr(\theta_i^1 = I | \lambda_i^0)$. After each subsequent meeting the players update their beliefs so $\mu_j^2 = E(\theta_i^2) = \Pr(\theta_i^2 = I | \lambda_i^0, \bar{\theta}_i^1)$. More generally, this can be expressed as;

$$\mu_j^t = E(\theta_i^t) = \Pr(\theta_i^t = I | \lambda_i^0, \{\bar{\theta}_i^{t-1}, \bar{\theta}_i^{t-2}, \dots, \bar{\theta}_i^{t-n}\}) \quad (3)$$

where n is the number of previous meetings (in the empirical section we set $n \leq 4$ for practical reasons).

Proposition 2. *Given that not all signals are empty (all priors are too low), two risk-averse players receive more non-empty signals on average than one risk-averse and one risk-loving player. Moreover, one risk-averse and one risk-loving player receive more non-empty signals*

on average than two risk-loving players. Hence, in general, risk-averse players update their beliefs more frequently than risk-loving players.

Proof. For a signal to be non-empty, a draw offer has to be made which is only made by type I players, hence, $\Pr(\bar{\theta} \neq \Theta | \theta_i = II, \theta_j = II)$ always equals zero. Since the initial belief of a player is a function of the game history of the opponent, i.e. $\mu_i^1(\lambda_j^0)$, a player sets a lower initial belief for a type II opponent than for a type I opponent. Given that not all priors are too low, it follows that $\Pr(\bar{\theta} \neq \Theta | \theta_i = I, \theta_j = I) > \Pr(\bar{\theta} \neq \Theta | \theta_i = I, \theta_j = II) > \Pr(\bar{\theta} \neq \Theta | \theta_i = II, \theta_j = II) = 0$. ■

4.2 Hypotheses and Implications of the Model

The way the game is constructed implies that the payoffs and equilibrium of the game are asymmetric for the two types. As has been discussed above, the best reply for a type II player is to avoid arranged draws. It follows that a type II player is never the first player to act and, consequently, this player is playing the best reply from the beginning. Hence, a type II player is not expected to update the beliefs and, therefore, there will be no learning for a type II player. Note that a risk-loving player will not always be a type II player so we still expect a risk-loving player to learn.

Our hypotheses in this paper are that; *i*) Chess players learn about their opponent's utility in repeated meetings and adapt their future strategies accordingly, *ii*) Risk-averse players learn faster than non-risk-averse players.

Let $\Delta\mu_i$ denote the learning effect (the updating of beliefs) for player i at time t , where $\Delta\mu_i = \mu_i^{t+1} - \mu_i^t$. Given that at least some priors are sufficiently high, we have five cases;

- a) $\Delta\mu_i | (\bar{\theta}_j^t = I, \theta_i^t = I) > 0$ [type I players learn and update their beliefs *upwards*]
- b) $\Delta\mu_i | (\bar{\theta}_j^t = I, \theta_i^t = II) = 0$ [type II players do not update]
- c) $\Delta\mu_i | (\bar{\theta}_j^t = \Theta) = 0$ [no update without new information]
- d) $\Delta\mu_i | (\bar{\theta}_j^t = II, \theta_i^t = I) < 0$ [type I players learn and update their beliefs *downwards*]
- e) $\Delta\mu_i | (\bar{\theta}_j^t = II, \theta_i^t = II)$ [not possible, no signal without a type I player]

In this paper we cannot distinguish between *rejected* draw offers and *no* draw offer, so we compare *accepted* draw offers to rejected or no draw offers (lumped together). Hence, we investigate empirically whether $\mu_i^{t+1} | (\theta_i^t = I)$ is greater than μ_i^t .¹⁴

If the prior $\mu < \frac{\gamma}{U^{draw} - EU + \gamma}$, the corresponding player will never offer a draw

(under ceteris paribus conditions). This implies that if the prior is sufficiently high, there will be a change of equilibrium and a divergence between the risk-averse and risk-loving equilibrium. However, if the prior is too low, there will be no change of equilibrium.¹⁵ Following from propositions 1 and 2, the model also predicts that more risk-averse players will learn more than risk-loving players as they update their beliefs more frequently.

Due to the simplicity and transparency of the game, the intelligence and cognitive ability of the players, together with the high level of information (game history and Elo rating), we find it reasonable to assume that the players, given their beliefs, are able to choose the optimal strategy. As the relative skill is known in every game, the players only have to know whether they prefer an arranged draw or not and evaluate whether they are willing to take the risk of offering a draw.¹⁶ Even if they were to fail playing best reply at every moment, they will probably be very close to optimal behavior.¹⁷

5 Data and Econometric Model

5.1 Data

The data in this study were obtained from ChessBase 10, a database collection with more than 1.5 million chess games played in high-level international chess events by expert chess players (Chessbase 10 has more than 5 million games in total). The resulting data set contains about 30,000 players from 140 countries. Two levels of data are available, player-specific information and game-specific information. The name, year of birth, nationality and gender of a player are available. For every game there are data on the names and Elo ratings

¹⁴ Here, players are supposed to take into account their chess-related payoff and their risk preferences but not, for instance, friendship or other kind of social preferences. If two friends, siblings or a couple play against each other, they may receive a higher utility from accepting an arranged draw although they are far superior. However, in this paper we want to exclude such effects and focus on genuine learning effects. Such potential confounders are discussed in section 6.2.

¹⁵ Such implications are common, see for instance Young (2004) and Fudenberg *et al.* (2004).

¹⁶ Recall that Moul and Nye (2009) found that the Soviet players were able to improve their tournament score by agreeing to arranged draws at the correct moment.

¹⁷ We suggest that the players play in accordance with the concept of *self-confirming equilibrium*. This equilibrium concept was defined by Fudenberg *et al.* (2004). This equilibrium concept has similarities with the Nash equilibrium, but here the player is not expected to know what has not yet been learnt.

of the two players, year of the game, number of moves and score. The years included in this study range from 1997 to 2007 and the minimum Elo rating required is 2000, above which players are considered to be experts. As regards the information on a player's nationality, we have grouped the countries in regions based on geographic lines and chess popularity. The regions with the highest number of chess players are Western Europe, Eastern Europe and the former Soviet Union. These three regions account for about 90 percent of the expert chess players in the world. Western Europe alone accounts for 53 percent, Eastern Europe for 24 percent and the former Soviet Union for about 13 percent. Latin America, North America¹⁸, Africa and Asia account for less than 10 percent.

5.2 Econometric Model

The unit of observation in the econometric model is the game. The dependent variable, Y , takes on the value 1 in the case of an arranged draw, and 0 otherwise.

We condition on a set of information for each player. This set of information is the same as that of the players, available from the database, as discussed in section 2.3. Information about the two individuals participating in a game is used to construct our explanatory variables. For each game, we have two sets of characteristics, one for white, X_w , and one for black, X_b . Elo rating (playing skill), gender, nationality, age polynomials, number of games, OR risk preferences, regional and year dummies are included as control variables. Included in X are also player-specific history variables, such as the tendency for each player to play safe and risky openings with the relevant color (the fraction of games with a given color against players above 2000 where the player played safe, or risky), the tendency for each player to end a game in arranged draws, the total number of past games, etc. Thus, we control explicitly for the historical risk profile, the opening preferences, and the ability of each player.

In addition, we include information on how often two players have met before with a set of indicators for the order of the game. Furthermore, we construct a set of 'lagged' dependent variables, taking the value 1 if there was an arranged draw in the previous game (and for 2, 3 and 4 time lags) between two players, 0 otherwise (i.e. including wins, losses and normal draws).

Thus, regressing the dependent variable, an indicator for an arranged draw, on the lagged dependent variable, a positive coefficient implies an increasing probability of the

¹⁸ As the U.S. Chess Federation applies a different rating of the playing strength many American players are missing in the data.

occurrence of an arranged draw with repeated meetings in general. Conditioning as well on all the information in the database, we argue that the parameters on the lagged dependent variables can be given a causal – learning – interpretation. Interacting the lagged dependent variables with the risk profile of each player, as measured by the fraction of previous games ended in arranged draws will show whether risk-averse players learn faster or not.

Each game is treated like an independent observation, except that we calculate robust standard errors that are cluster at the player level. As regards the lagged variables, the current encounter is the dependent variable, whereas the first, second, third and fourth previous encounters correspond to the four lagged dependent variables. In cases where two players meet more than five times, the fourth lag is 1, lending it an interpretation of all learning taking place at the fifth game and later. The model is estimated using a linear probability model (OLS).

6 Results

6.1 Empirical Results

The main results from the estimations are presented in Table 1. Due to the complexity to overview the results, we present the learning effect coefficients in Figure 3 which also displays how the learning effects vary across risk level.

From Table 1 we see that the ‘Opening Repertoire (OR) risk-averse’ coefficients are as expected, i.e. they are positive and the probability of an arranged draw increases even more if both players are OR risk-averse. We also see that players of opposite sex have a smaller probability of an arranged draw. Being of the same gender has a positive impact on the probability of an arranged draw. Moreover, the coefficient for having the same nationality is positive. It is possible that players of the same gender or nationality have a higher prior belief which would lead to faster learning. This may also capture part of a friendship effect, see section 6.2.

Due to the complexity of the model, especially with respect to the marginal effect of past games on the likelihood of the present outcome, we present the marginal learning effect of past games in Figure 3

Table 1. Estimation results for the main regression (LPM estimated with OLS).

	Dep var: AD
Opening Repertoire (OR) risk-averse, white	.0111 (.0009)***
OR risk-averse, black	.0192 (.0011)***
Both OR risk-averse	.0141 (.0014)***
Female, white	-.0040 (.0018)**
Female, black	-.0078 (.0015)***
Both female	.0162 (.0026)***
Same nationality	.0229 (.0008)***
<i>Learning effects</i>	
Arranged draw (AD) 1 lag	.0582 (.0128)***
AD 2 lags	.0510 (.0186)***
AD 3 lags	-.0272 (.0228)
AD 4 lags	.0970 (.0281)***
AD 1 & 2 lags	.0363 (.0141)***
AD 1, 2 & 3 lags	.06270 (.0273)**
AD 1, 2, 3 & 4 lags	.0221 (.0345)
AD 1 lag * AD risk (white)	.3184 (.0820)***
AD 2 lag * AD risk (white)	.1760 (.1217)
AD 3 lag * AD risk (white)	.8851 (.1572)
AD 4 lag * AD risk (white)	-.2521 (.1719)
AD 1 lag * AD risk (black)	.2313 (.0758)***
AD 2 lag * AD risk (black)	.1764 (.1242)
AD 3 lag * AD risk (black)	.4421 (.1582)***
AD 4 lag * AD risk (black)	-.1499 (.1919)
AD 1 lag * AD risk (white & black)	-.0345 (.4326)
AD 2 lag * AD risk (white & black)	-.4551 (.6719)
AD 3 lag * AD risk (white & black)	-3.8259 (.9164)***
AD 4 lag * AD risk (white & black)	1.6205 (1.0094)
Normal draw 1 lag (≥ 20 moves)	.0262 (.0020)***
Normal draw 2 lags (≥ 20 moves)	.0178 (.0035)***
Normal draw 1 & 2 lags	-.0075 (.0056)
OR risk-averse share, white player	-.0068 (.0028)**
OR risk-loving share, player	-.0077 (.0035)**
AD risk, player	.4398 (.0078)***
Number of historic games, player	-.00004 (.000004)***
OR risk-averse share opponent	-.0200 (.0024)***
OR risk-loving share opponent	-.0165 (.0030)***
AD risk, opponent	.4089 (.0061)***
Number of historic games, opponent	.000003 (.000004)
2 nd encounter	.0052 (.0013)***
3 rd encounter	.0014 (.0025)
4 th encounter	-.0051 (.0035)
5 th encounter	-.0216 (.0032)***
Elo, (white & black)	Yes
Age, Age-2 (white & black)	Yes
Regional dum. (white & black)	Yes
Year dummies (date of game)	Yes
Constant	-.1821 (.0104)***
Number of games/players	744,307 / 32,093

Notes: Western Europe is used as a reference group for the regional dummies. Robust standard errors in parentheses, clustered at player level. * significant at 10%; ** significant at 5%; *** significant at 1%. AD is short for Arranged Draw, and OR is short for Opening Repertoire. 'White' refers to the player holding the white pieces, 'black' refers to the player holding the black pieces.

Figure 3 also displays how the learning effect varies across risk preferences, where the 10th percentile is the 10th percentile in the distribution of the variable measuring the fraction of past games for a given player that ended in an arranged draw. The 10th percentile thus represents the risk-loving chess players, while the 90th percentile represents those that are most risk-averse. Figure 3 shows that in general, learning about each other's preferences occurs, reflected in the positive marginal effect of past arranged draws on the probability of another arranged draw. Moreover, as predicted by the theoretical model, there is more (and faster) learning for more risk-averse players. For the 90th percentile we find that the probability of an arranged draw is 55 percentage points greater when the four previous meetings ended in arranged draws than when no arranged draw has been agreed to in the past. For the 10th percentile this effect is 30 percentage points.¹⁹

Figure 3. The learning effects for 1-4 earlier AD:s for different levels of risk preferences, where the 10th percentile is the most risk-loving and the 90th the most risk-averse group.

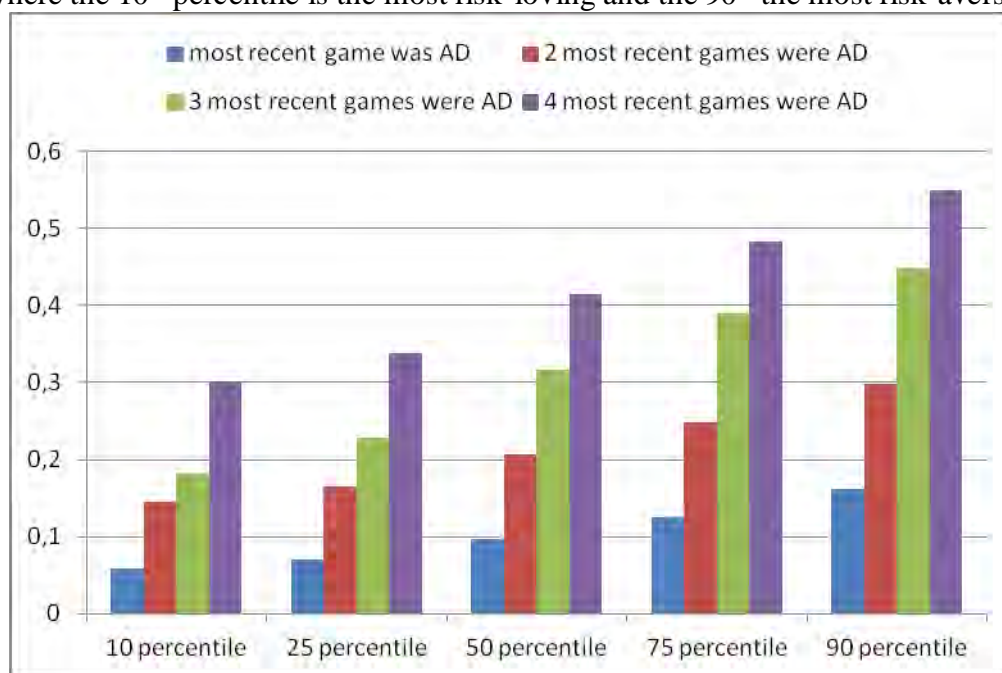
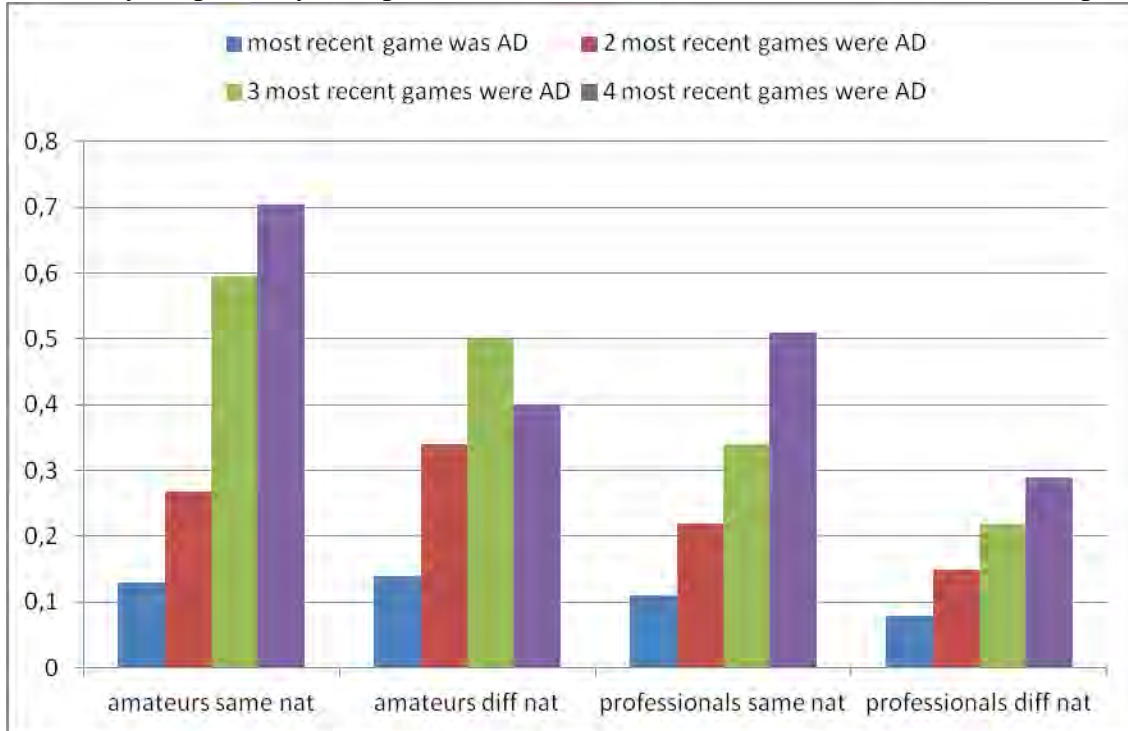


Figure 4 displays the results for amateurs and professionals, and same-nationality, different-nationality games, separately. The purpose is to see if the pattern remains for professional players, especially when of different nationalities. It is less likely that friendship has an

¹⁹ The effect on the probability of arranged draws from having played normal draws (twenty moves or more) in earlier encounters, see Table 1, is much smaller than the effect from having played earlier arranged draws. The effect is about 2 to 3 percentage points. Having played normal draws in earlier meetings might increase the probability of an arranged draw as it may suggest that a draw is the most probable outcome in any case. There is, however, limited evidence for this type of learning.

important impact on the probability of arranged draws among professional players and should also be less common across nationalities.²⁰

Figure 4 – learning effects for different subgroups, amateurs with same and different nationality, respectively, and professionals with same and different nationalities, respectively.



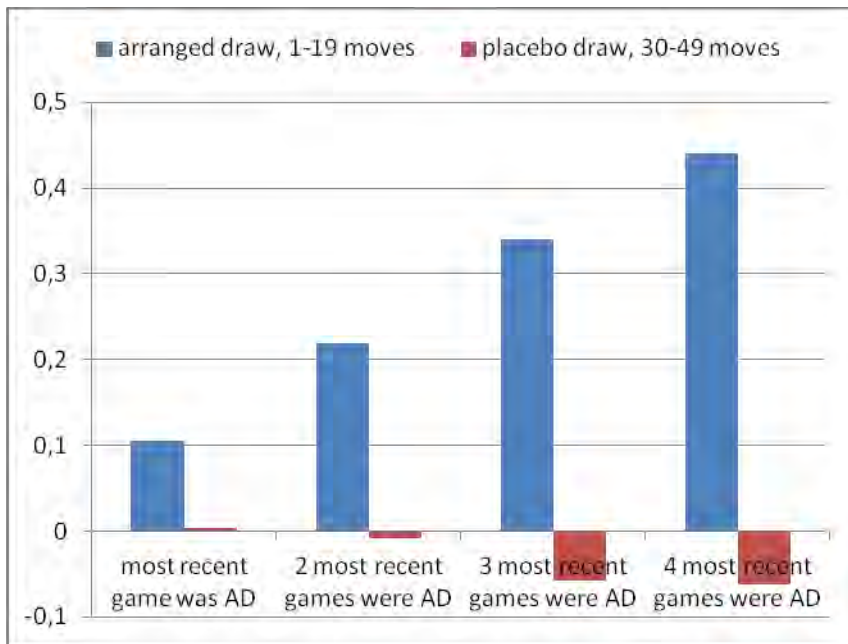
Note: The marginal effects are evaluated at the median of the AD risk aversion distribution.

The fact that the learning effect also appears for professional players and for players of different nationalities reduces the probability of friendship driving the results. There is no significant difference between same and different nationality subgroups. There is, however, some difference in the size of the effect between amateurs and professionals. This may be due to the fact that professional players typically have more a priori information about the opponents than the amateurs have.

As a placebo analysis, Figure 5 shows the results from a regression, where arranged draws, defined as draws in fewer than twenty moves, are replaced as the outcome variable by draws in 30-49 moves. As is seen in Figure 5, the effect from the placebo ‘treatment’ is close to zero (actually slightly negative).

²⁰ Note that the findings by Moul and Nye (2009) that the Soviet players tended to collude, relate to the absolute world elite, and hold only in countries where there was a strong political pressure on the players to perform. Players at lower levels or in countries with greater political freedom are not very likely to ‘sacrifice’ themselves for the nation in individual tournaments.

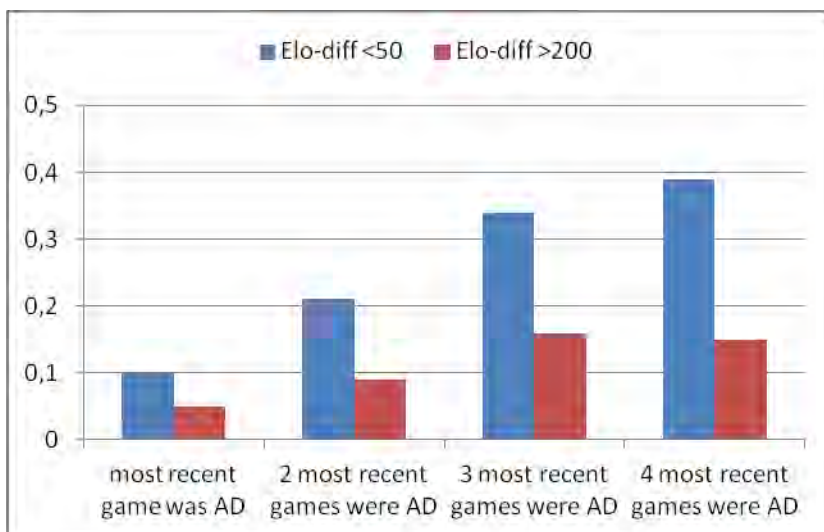
Figure 5 - Placebo (draw in 30-49 moves) and real learning effects (draw in 1-19 moves).



Note: The marginal effects are evaluated at the median of the AD risk aversion distribution.

Figure 6 presents the results when the difference in playing skill, the Elo difference, is smaller than 50 Elo points and larger than 200 Elo points. If a player is substantially more skilled than the opponent, her belief about the opponent's drawing preferences will be close to one while the corresponding belief of the opponent will be close to zero. For this reason, when the difference in playing skill is large between two players, the learning is expected to develop slower since the superior player's belief cannot increase (much) more and the inferior player will suspect that the opponent's preference for an arranged draw in the past was temporary.

Figure 6 – learning effects when the Elo difference is < 50, and > 200 points, respectively.



Note: The marginal effects are evaluated at the median of the AD risk aversion distribution.

As predicted by the theoretical model, Figure 6 shows faster learning when the Elo difference is small.²¹

6.2 Potential Confounders

One possible confounder in this context is the existence of friendship (friends, siblings or similar) between two players, which could increase the non-economic payoff from an arranged draw. In general, it is extremely difficult to obtain information about such payoffs but we can control implicitly for several factors which makes it less likely that the results are driven by friendship.

First, the effects found in this paper are large and monotonically increasing. If these effects are driven by friendship, there would have to be a substantial, monotonic increase in friendship from *one game to another*. It seems more plausible that the effects are due to updating of the beliefs. Second, we find that risk-averse players learn about twice as fast which, if the results are driven by friendship rather than real learning effects, would indicate that risk-averse players have a more developed friendship. Although possible, there is no consensus pointing in this direction. Third, we find learning effects also for professional players. It seems less likely that people playing chess for a living should let friendship interfere with “business”. Fourth, we find no learning effects when replacing arranged draws with the “placebo” outcome, i.e., draws in 30-49 moves. Since draws, regardless of the number of moves that have been played when agreed to, give half a point, friendship should affect the outcome also in the placebo regression.²² Fifth, we also find that the learning effects are lower when the Elo difference is larger. This is in line with the learning effects predicted by the theoretical model but cannot be considered to be a typical pattern if driven by friendship. If anything, the pattern should be the reversed as rivalry usually increases when players are more equal in skills.

Finally, we should point out that we find some support for players of similar gender or nationality showing faster learning effects, see Table 1. One can argue that friendship is likely to be stronger within the same gender or nationality than otherwise. This argument would give some support to friendship being the driving factor behind the results. However, we

²¹ When comparing male and female players, there is, although not significantly different from each other, a small tendency for female players to learn faster than male players. This is in line with the consensus that women are more risk-averse than men.

²² The effects from including normal draws (twenty moves or more) as explanatory variables are very small compared to arranged draws.

propose an alternative interpretation. It is possible, and perhaps also likely, that the initial prior is higher when the players have the same nationality or gender, i.e., it may be easier to interpret pre-game signals when having a similar cultural background.²³

7 Conclusions

We conclude that there is rational learning involved in repeated chess games between two chess players. Past outcomes between two players affect the beliefs about the opponent's preferences in future games, although the general type of the opponent is known in advance. The magnitude of the effects is large with roughly an additional 10 percentage points higher probability for each additional previous arranged draw. The theoretical model we develop for this setting, predicts that risk-averse players learn faster than risk-loving players, as they receive more signals. We find strong support for this being the case. The learning effect is about twice as large for the most risk-averse players (at the 90th risk percentile) compared to the learning effects for the most risk-loving players (at the 10th risk percentile). The fact that risk-averse players learn faster depends on the construction of this particular game. The conclusion from this is rather that people, although they have the same intellectual and cognitive ability, may learn differently depending on their preferences, in this case risk preferences.

The learning effects survive several sensitivity tests. We carry out a “placebo treatment” which, as predicted, shows no learning effects at all. Furthermore, large rational learning is found both for professionals and amateurs and we also find that the learning effects decrease when the difference in playing skill increases. The fact that the results are insensitive to such tests strengthens the liability of the findings, as it is less likely that they are driven by friendship or other “social” preferences.

Studying intelligent and productive expert chess players with high cognitive capacity that are well equipped with adequate data and specialized computer programs is important when analyzing rational learning. By considering such an extreme group, we may obtain a boundary result as regards rational learning. Although not representative for the society as a whole the findings resulting from analyses on this subgroup may supply researchers with an upper bound which can be used as a benchmark in future research on these topics.

²³ Such arguments have been suggested by, for instance, Cornell and Welch (1996) and Lang (1986).

Ultimately, with this study we hope to demonstrate how useful large-scale chess data from international expert chess games can be for economic research. As a chess game constitutes a highly controlled environment with homogenous rules across countries, it can be seen as one of the largest registered field experiments on economic behavior.

References

- Camerer, Colin, (2003), *Behavioral Game Theory: experiments in strategic interaction*. Princeton University Press, Princeton, NJ.
- Chabris, Christopher F. and Mark E. Glickman, (2006), “Sex Differences in Intellectual Performance – Analysis of a Large Cohort of Competitive Chess Players,” *Psychological Science*, 17, pp. 1040-1046.
- Charness, Neil, (1992), “The impact of chess research on cognitive science,” *Psychological Research*, 54, pp. 4-9.
- Cornell, Bradford, and Ivo Welch, (1996), “Culture, Information and Screening Discrimination,” *Journal of Political Economy*, 104, pp. 542-571.
- de Firmian, Nick, (2009), *Batsford’s Modern Chess Openings*, Batsford, London.
- Elo, Arpad E., (1978), *The Rating of Chess Players, Past and Present*, Batsford, London.
- Falk, Armin and James J. Heckman, (2009), “Lab Experiments Are a Major Source of Knowledge in the Social Sciences,” *Science*, 326, pp. 535-538.
- Fudenberg, Drew, Eddie Dekel and David K. Levine, (2004), “Learning to Play Bayesian Games,” *Games and Economic Behavior*, 46, pp. 282-303.
- Gobet, Fernand, (2005), “Chunking Models of Expertise: Implications for Education,” *Applied Cognitive Psychology*, 14, pp. 183-204.
- Lang, Kevin, (1986), “A Language Theory of Discrimination”, *Quarterly Journal of Economics*, 101, pp. 363-382.
- Levitt, Steven and John List, (2007), “What do laboratory experiments tell us about the real world,” *Journal of Economic Perspectives*, 21, pp. 153-174.
- Levitt, Steven, John List and Sally Sadoff, (2009), “Checkmate: Exploring Backward Induction Among Chess Players,” NBER Working Paper Series, <http://www.nber.org/papers/w15610>
- Moul, Charles C. and John V. Nye, (2009), “Did the Soviets Collude? A Statistical Analysis of Championship Chess 1940–1978,” *Journal of Economic Behavior & Organization*, 70, pp.10-21.

Palacios-Huerta, Ignacio and Oscar Volij, (2009), “Field Centipedes,” *American Economic Review*, 99, pp.1019-1035.

Roring, Roy W., (2008), “Reviewing expert chess performance: a production based theory of chess skill,” Dissertation, Department of Psychology, Florida State University.

Ross, Philip E., (2006), “The Expert Mind,” *Scientific American*, 295, pp. 64-71.

Ross, Daniel, (2007), “Arpad Elo and the Elo Rating System”

<http://www.chessbase.com/newsdetail.asp?newsid=4326> .

Simon, Herbert A., (1955), “A Behavioral Model of Rational Choice,” *Quarterly Journal of Economics*, 69, pp. 99–118.

Young, Peyton, (2004), *Strategic learning and its limits*, Oxford University Press, Oxford and New York

Appendix

The Elo Rating System²⁴

The following description aims to show how winning probabilities are calculated and how Elo rates of chess players are subsequently updated.

In the late 1950s, Arpad Elo, a physicist and a devoted chess player, introduced a new system of classifying the strength of chess players. By observing results from chess tournaments, he noted that the distribution of individual performances in chess resembled a normal distribution. On the basis of his observations, he introduced a point scale, where he determined the standard deviation to be 200. The distribution relates to the difference in ability between two players. Defining μ as the difference in Elo strength between two players, this gives us the following probability density function (pdf), i.e. the marginal probability of winning:

$$f(x) = \frac{1}{\sigma' \sqrt{2\pi}} e^{-\frac{(\mu)^2}{2\sigma'^2}} \quad (1)$$

As there are two participants, each of them with an assumed performance deviation of 200 Elo points, the standard deviation used in (1) can be rewritten as follows:

$$\sigma' = \sqrt{\sigma_1^2 + \sigma_2^2} = \sqrt{2\sigma^2} = \sigma\sqrt{2} = 200\sqrt{2} \quad (2)$$

²⁴ The following section is to some extent drawn from Ross (2007).

$$\int \frac{1}{\sigma'\sqrt{2\pi}} e^{\frac{-(\mu)^2}{2\sigma'^2}} d\mu = \frac{1}{400\sqrt{\pi}} \int e^{\frac{-(\mu)^2}{400^2}} d\mu \quad (3)$$

For example, this is saying that the probability of winning is 76 [24] percent if one player has 200 Elo points more [less] than his/her opponent. When two players are equally strong (i.e. an Elo difference equal to zero), the most likely outcome is a draw.

Elo ratings: A sequential estimate of strength

The probability of winning as shown in equation (3) is used to update a player's Elo rating. The algorithm for this reads as follows:

$$Elo_new = Elo_old + (Score - Prob(\text{winning})) * k \quad (4)$$

Here Elo_old is the Elo rating before the game starts, whereas Elo_new is the updated rating. The $Score$ indicates the actual outcome of a game, where a win [loss] is valued as 1 [0], and a draw counts as .5 point. The coefficient k is a weighting factor that determines how much the outcome of a game counts for a player's Elo rating. It is determined by the number of games played, i.e. the less experienced a player is, the higher the k .

Example: Let us hypothesize that there are two players with an Elo difference of 100 Elo points. This corresponds to an expected score of approximately 62.5% (or 5/8) for the superior player and 37.5 (or 3/8) for the inferior player. The players will then choose the set of probability parameters (p and q) as to maximize their utility. For instance, the *inferior* player could choose between two different parameter combinations, in the first $p=2/8$ and $q=2/8$ whereas in the second $p=1/8$ and $q=4/8$. Both strategy choices render the same expected score, 3/8. A more risk-averse choice is $\frac{2}{8}U(1) + \frac{2}{8}U(1/2) + \frac{4}{8}U(0)$ while $\frac{1}{8}U(1) + \frac{4}{8}U(1/2) + \frac{3}{8}U(0)$ is more risk-loving. Since $[\phi - (1 - \phi)] = .375 - .625 = -.25$ or $\frac{3}{8} - \frac{5}{8} = -\frac{2}{8}$, we have $p = r - .25$, i.e. the probability of a win is lower than the probability of a loss due to the fact that the player is inferior by 100 Elo points. In this example, $p \in [0, \frac{3}{8}]$, $q \in [0, \frac{6}{8}]$ and $r \in [\frac{2}{8}, \frac{5}{8}]$. The second strategy is more risk-averse since it leads to a higher probability of a draw.