

How do adults reason about their opponent? Typologies of players in a turn-taking game

Tamoghna Halder (thaldera@gmail.com)
Indian Statistical Institute, Kolkata, India

Khyati Sharma (khyati.sharma27@gmail.com)
Indian Statistical Institute, Kolkata, India

Sujata Ghosh (sujata@isichennai.res.in)
Indian Statistical Institute, Chennai, India

Rineke Verbrugge (l.c.verbrugge@rug.nl)
Institute of Artificial Intelligence, University of Groningen, The Netherlands

Abstract

This paper reports a construction of typologies of players based on their strategic reasoning in turn-taking games. Classifications have been done based on latent class analysis and according to different orders of theory of mind, and exploratory validations have been provided for the resulting classifications. Finally, interaction of the typologies described by these classifications is discussed towards achieving a common perspective of typologies of players originating from various aspects of strategic thinking.

Keywords: social cognition; higher-order theory of mind; strategic games; turn-taking games

Introduction

Theory of mind (ToM) is the ability to attribute beliefs, desires, and intentions to other people, in order to explain, predict and influence their behavior. Even though ToM has been widely studied in the cognitive sciences, relatively little research has concentrated on people's reasoning about their opponents in turn-taking games. We speak of zero-order reasoning in ToM when a person reasons about world facts, as in "Anwasha wrote a novel under pseudonym". In first-order ToM reasoning, a person attributes a simple belief, desire, or intention to someone else, for example in "Khyati *knows* that Anwasha wrote a novel under pseudonym". Finally, in second-order ToM reasoning, people attribute to others mental states about mental states, as in "Khyati *knows* that Soumya *thinks* that Anwasha did not write a novel under pseudonym".

One way of studying the cognitive basis of ToM in a controlled experimental setting is the use of turn-taking games. By investigating the underlying strategies used during these games, one can shed light upon the underlying cognitive processes involved—including ToM reasoning. In recent times, higher-order theory of mind has been the central focus of a lot of research papers that are based on experiments with games (see, for example, Camerer, 2003). Higher-order ToM reasoning also became an attractive topic

for logical analysis. These logical investigations often take recourse to game theory (Osborne and Rubinstein, 1994). In recent years, game-theoretic experiments have formed the backbone of research in behavioral game theory. In general, experimental studies are essential in studying social phenomena that govern and are governed by individual or collective human behavior.

The main focus of the current paper is to investigate which rules govern human strategic thinking, in order to develop a typology of players based on their cognitive strategies. This paper is based on experimental studies using turn-taking games. Such games are ubiquitous in our daily life – debates and deliberations, negotiations, coalition formation, and others. The marble drop games that we use, which are game-theoretic equivalents to Hedden and Zhang's (2002) 'matrix games', have been designed by Meijering. They are so-called perfect information games, in contrast with games like poker and bridge, in which players cannot see the others' cards (Osborne and Rubinstein, 1994). Such marble drop games have been used extensively to study various cognitive phenomena, especially those involving higher-order theory of mind (Meijering, van Rijn, Taatgen & Verbrugge, 2012; Ghosh, Meijering & Verbrugge, 2014; Meijering, Taatgen, van Rijn & Verbrugge, 2014). However, as far as we know, studies on the underlying typology of players in turn-taking games are very scarce. The questions arise like what kind of typology of players we are looking for and how a typology could be beneficial in studying the cognitive phenomena involved in playing turn-taking games. The idea is to come up with a list of basic properties or concepts that differentiate human strategic reasoners in terms of their reasoning approaches. People's approaches depend, among other factors, on the attributes that they assign to their opponents. Therefore, participants' analysis of their opponent's approach becomes an important factor in formulating the properties that we can apply in a useful typology, and thus ToM provides an essential concept for study in these cases.

In the literature on behavioral game theory, there is a natural tendency to analyze mostly the choices made by players at different turns of the game, thereby ignoring the data on how much time they have taken to make that choice, namely, the ‘response time’ data. Rubinstein (2014) does argue for the importance of response times and takes that data into account while discussing a typology of players in different games. Also, he discusses typologies that are beyond the traditional psychometric typologies originating from ‘type theory’ and ‘trait theory’ (Bateman, Lowenhaupt & Nacke, 2011). Rubinstein views the analysis from a game-theoretic point of view and therefore the use of robust statistical methods comes into the picture only for the sake of validating the game-theoretic implications. The current paper, in contrast, looks at an experiment from a statistical angle as well as from the viewpoint of theory of mind. Instead of defining typologies on the basis of game-theoretic approaches, we will consider the data and use statistical analysis to develop a new kind of domain-specific typology. Furthermore, to cross-validate the plausibility of the developed typology, the interplay of the developed typology with various degrees of rationality arising from theory of mind (namely zero-order, first-order and second-order theory of mind) will also be investigated. Thus, the focus of this paper is two-fold: to study the typology of players from the domain-specific viewpoint and to connect the gap between discretely originated player types. Finally, the study of such a typology of players helps to explain the differences between people’s cognitive attitudes when reasoning strategically and to better understand people’s possible behaviors in interactive situations. This in turn helps in deciding the controlling factors of people’s strategic reasoning processes, which can be used for modeling purposes in various disciplines, for example, economics, artificial intelligence, logic, and linguistics, where formal, behavioral and algorithmic studies of social phenomena are taken up. In this paper, the subsequent sections will focus on the structure of the experiment and associated data, classification based on latent class analysis, classification based on theory of mind, and the interaction of these two classifications for the purpose of exploratory validation. The Discussion presents conclusions and future directions.

Methods and Data

We provide a brief summary of the experimental games and the experimental procedure underlying the current work. The experiment was conducted at the Institute of Artificial Intelligence (ALICE) at the University of Groningen, The Netherlands. The games that were used in the experiment are given in Figures 1 and 2. In these two-player games, the players play alternately, therefore they are called *turn-taking* games. Let C denote the computer and P the participant. In the first four games (Figure 1), the computer plays first, followed by the participant. The players control two decision nodes each. In the last two games (Figure 2), which are truncated versions of two of the games of Figure 1, the participant moves first.

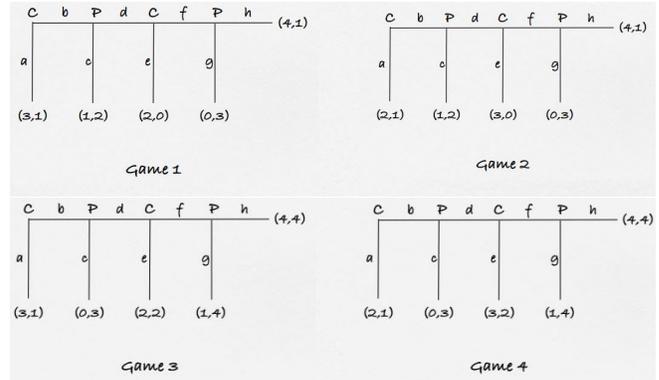


Figure 1: Structures of the games 1, 2, 3, and 4. The computer (C) plays first. The ordered pairs at the leaves represent pay-offs for the computer (C) and the participant (P), respectively.

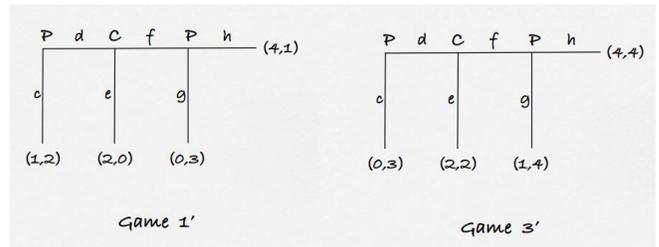


Figure 2: Structures of the truncated games 1' and 3'. The participant (P) plays first.

Before going any further, let us first explain two relevant ways of playing these games as prescribed by game theory—the backward induction (BI) procedure (Osborne & Rubinstein, 1994) and the extensive form rationalizability (EFR) concept (Pearce, 1984). See (Ghosh, Heifetz & Verbrugge, 2015) for a precise game-theoretical explanation of BI and EFR reasoning for the six experimental games presented in Figures 1 and 2. Informally, EFR takes into account an opponent’s past moves in order to assess that opponent’s future behavior, whereas BI only considers the opponent’s future choices and beliefs, and ignores the opponent’s past choices (“let bygones be bygones”). The question here is how the participant would play if her first decision node was reached; in games 1, 2, 3, 4, reaching the first P-node would already indicate that the opponent C had not opted for its rational decision, namely to go down immediately. Would the participant’s (P’s) decision depend on her opponent’s previous choice? Here, she would have to choose between continuing the game (by moving to the right, action *d*) and opting out (by moving down, action *c*).

According to the EFR concept, the expected behavior of the players would be as follows: *d* (instead of *c*) would be played more often in game 3 than in game 4, more often in game 1 than in game 2, more often in game 1 than in game 1', and more often in game 3 than in game 3'. The reason for taking EFR as our predictive concept rather than the

more popular BI concept is the fact that there have been a lot experimental validations (for example, Rosenthal, 1981) that show that people do not follow BI behavior in such turn-taking games of perfect information.

A group of 50 Bachelor and Master's students from different disciplines at the University of Groningen took part in the experiment. The participants played finite perfect-information games that were game-theoretically equivalent to the games depicted in Figures 1 and 2. However, the presentation was made such that participants were able to understand the games quickly, see Figure 3. In each game, a marble was about to drop. Both the participant and the computer determined its path by controlling the trapdoors: The participant controlled the orange trapdoors, and the computer the blue ones. The participant's goal was that the marble should drop into the bin with as many orange marbles as possible. The computer's goal was that the marble should drop into the bin with as many blue marbles as possible. In Figure 3, corresponding to game 1 of Figure 1, if the computer uses BI, it opens the top left blue trapdoor, leading to 3 blue marbles (its rational choice for this game). For, if the computer had opened the right blue trapdoor, the participant (also applying BI) would then have opened the right orange trapdoor to obtain 2 orange marbles, because had he opened the left orange trapdoor, the computer at its next move would have opened the left blue trapdoor, leaving the participant with 0 marbles; the computer would have done this, because otherwise at his bottom-most orange trapdoor, the participant would open the left orange trapdoor to attain 3 orange marbles, leading to no marbles for the computer.

In the experiment, however, the computer often makes an apparently irrational first choice, operationalized as follows. For each game item, the computer opponent had been programmed to play according to plans that were best responses to some plan of the participant. This was told to the participants in order to bring them all on a uniform level with respect to pre-knowledge of the game and to ensure that their behavior is independent of their exposure to computer games in personal life. In fact, each participant at first played 14 practice games with which they got a feel of the games before the start of the experiment. In the actual experiment, they played 48 games divided in 8 rounds, each comprised of 6 different game structures that were described above (see Figures 1 and 2). Different graphical representations of the same game were used in different rounds. A break of 5 minutes was given after the participant finished playing 4 rounds of the experimental games. At some points during the experimental phase, the participants were asked a multiple-choice question, as follows:

“When you made your initial choice, what did you think the computer was about to do next?”

Three options were given to the participants, regarding what they thought to be the likely next choice of the computer:

- *I thought the computer would most likely play e.*
- *I thought the computer would most likely play f.*
- *Neither of the above.*

In addition to the basic information on age, gender and departmental affiliation of the participants, for each game, for each round of the game, we collected the following data:

- *Participant's decision at his/her first decision node*, if the node was reached. In particular, whether move *c* or *d* had been played (cf. Figures 1 and 2);
- *Participant's decision at his/her second decision node*, if the node was reached. In particular, whether move *g* or *h* had been played (cf. Figures 1 and 2).

Moreover, for each participant, we collected the following data:

- *Participant's answer* to the above-mentioned multiple-choice question at the end of the rounds in which it was asked. In particular, whether the answer was *e* or *f* or undecided was noted.
- *Answering Time*: Time taken by the participant in giving the answer, i.e. the time between the moment the question appeared on the screen and the moment he/she clicked on his/her choice of answer.

Ghosh, Heifetz and Verbrugge (2015) show that overall, participants do pertain to EFR behavior in many of the games, even though in some cases there can be more mundane explanations for their choices. In the current paper, we try to get a more precise sense of how participants are reasoning, by distinguishing several types of players.

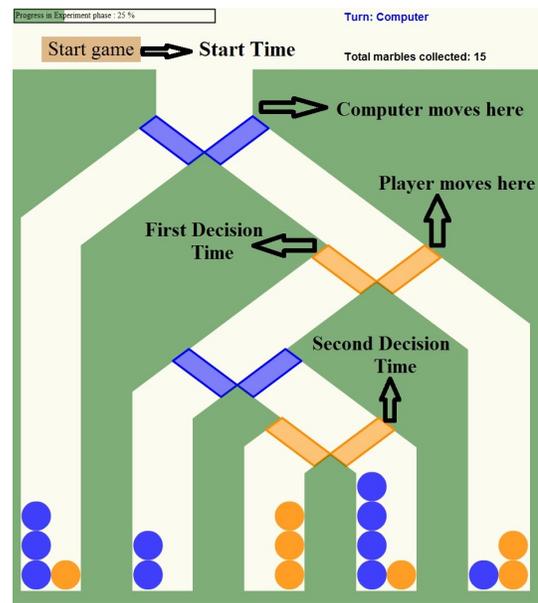


Figure 3: Graphical interface of an example game item.

Results: Latent Class Analysis

Latent class analysis (LCA) is a statistical method that can be applied to classify binary, discrete or continuous data in a manner that does not assign subjects to classes absolutely,

but with a certain probability of membership for each class (Goodman, 1974). LCA can be used to explore how participants can best be distinguished according to reasoning strategies, in cases where no fixed set of reasoning strategies has been defined in advance. Raijmakers, Mandell, van Es and Counihan (2014) have profitably applied LCA to the analysis of children’s reasoning strategies in turn-taking games. For the current experiment, the participants were categorized into certain classes based on their choices, c or d , at the first decision node in the game items corresponding to games 1, 2, 3 and 4 of Figure 1. The LCA was performed using the statistical software R, with 25 estimated parameters and 25 residual degrees of freedom.

The data for 50 participants were separated into two sets: the set containing the first three rounds and the set containing the last three rounds for each game; in each of the six rounds, the first decision node of a participant was reached. The participants were classified into two groups based on their behavior in each set of three rounds. Figures 4.1 and 4.2 show the graphs depicting their choices of ‘ c ’ in each of the rounds in each of the games (g_{ij} denotes behavior at the j^{th} round of the i^{th} game).

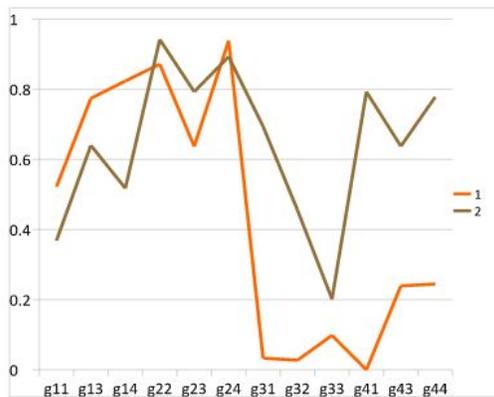


Figure 4.1: Graphical representation of LCA for the set containing the first three rounds for each game.

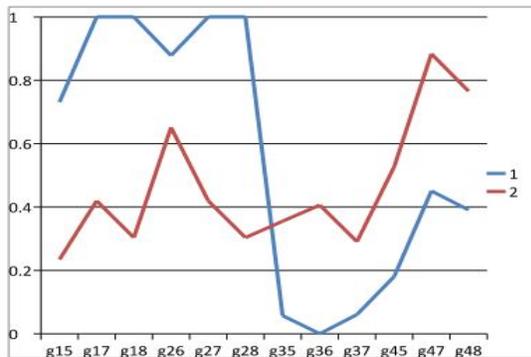


Figure 4.2: Graphical representation of LCA for the set containing the last three rounds for each game. The different predicted groups are denoted by different colors in Figures

4.1 and 4.2. Evidently, group 1 behaved in an expected fashion (akin to EFR behavior) in both cases, compared to the more random behavior of the other group. Considering group 1 for both sets of rounds, 24 common participants were noted down, who were predicted to behave in an expected fashion in all the rounds. The groups that resulted from the latent class analysis are as follows:

- a) Group 1: Playing in an expected fashion in both the initial three rounds and the later three rounds; there were 24 such players.
- b) Group 2: Not playing in an expected fashion in the initial three rounds but playing in an expected fashion in the later three rounds; there were 9 such players.
- c) Group 3: Playing in an expected fashion in the initial three rounds but not playing in an expected fashion in the later three rounds; there were 7 such players.
- d) Group 4: Not playing in an expected fashion in either the earlier or the later set of three rounds; there were 10 such players.

Statistical Typologies

On the basis of the above analysis, we propose the following statistically developed typology of players:

- 1) *Expected*: the 24 players who belong to group 1 above;
- 2) *Learner*: the 9 players from group 2 above;
- 3) *Random*: the 17 players from groups 3 and 4 combined.

For further statistical validations of the proposed typologies, we tested a number of hypotheses using standard statistical methods. One such hypothesis is to check whether the answering time is more in case of *expected* players than *random* players. The intuition behind this hypothesis is that a person who is playing in an expected fashion or learning to do so is bound to answer more “sensibly” and therefore would pay greater attention in choosing a correct option than a person who is playing less sensibly (random), cf. Rubinstein (2014). This hypothesis was tested twice using two sample t-test for difference of means, firstly *Expected* versus *Random* and secondly *Expected+Learner* versus *Random*. In both cases, our null hypothesis of equality of means was rejected at 5% level of significance (p-values 0.02 and 0.04, respectively). Hence, we may regard that the *Expected* and *Learner* players took more time in answering than the players termed as *Random*.

As a conclusion of the above analysis, we can regard that statistically developed typologies proposed above are robust at 5% level of significance.

Further Exploratory Validations

Each participant was asked the multiple-choice question about the most likely behavior of the computer opponent a

number of times (see the Methods section). We noted whether their answers were correct (that is, whether the answers corresponded with their actions before), corresponding to each round at which they were asked the question. Table 2 shows the exploratory results. The table justifies the way these groups were considered, since the number of wrong answers per player for the different groups hints at the fact that random moves led to more mistakes in answering the multiple-choice question.

Table 2: Table of summary for the wrong answers according to the latent class analysis

Group	Number of participants	Total number of wrong answers	Number of persons who gave all correct answers
1	24	28	7
2	9	17	3
3	7	34	0
4	10	17	3

Results: Theory of Mind Study

At the completion of the game-theoretic experiment, each participant was asked to answer the following final question:

“When you made your choices in these games, what did you think about the ways the computer would move when it was about to play next?”

The participant needed to describe in his or her own words, the plan he or she thought was followed by the computer on its next move after the participant’s initial choice. Based on their answer, 48 players were classified into three types according to the order of theory of mind exhibited in their answer to the final question. These were the types:

- a) *Zero-order players*, who did not mention mental states in their answer; there were 5 such players.
- b) *First-order players*, who presented first-order theory of mind in their answer; there were 27 such players;
- c) *Second-order players*, who presented second-order theory of mind in their answer; there were 16 such players.

This classification, as mentioned above, was done by manual scrutiny of each answer. Typical answers from each group are as follows:

- a) Zero-order answers: “It would repeat its former choice in the same situation.”
- b) First-order answers: “I thought the computer took the option with the highest expected value. So if on one side you had a 4 blue + 1 blue marble and on the other side 2 blue marbles he would take the option $4+1=2.5$.”

- c) Second-order answers: “...I thought the computer anticipated that I (his opponent) would go for the bin with the most orange marbles in his decision to open doors. This could lead to him getting less marbles than ‘expected’ because I would choose a safe option (3 marbles) over a chance between 4 marbles or 1 (depending on the computer’s doors).”

Statistical Validation

Based on the above three types of players (i.e. zero-order, first-order, and second-order players), we set up different hypotheses. Intuitively, one can expect that the players adopting second-order theory of mind would take maximum time to make a decision at the first decision node in comparison to players adopting first-order theory of mind and that people adopting zero-order theory of mind would take the least time among all three classes. This fact was validated statistically by performing difference of means test on the response time data of the first decision node for the three classes. We tested the hypotheses at 5% level of significance.

Combining the results, we found that $\mu_s > \mu_f > \mu_z$ for first decision time. Here, μ_s stands for the mean first decision time of second-order players, μ_f and μ_z denotes the same for the first-order and zero-order players, respectively.

Reviewing the results obtained, we can conclude that typologies based on theory of mind are statistically valid and robust at 5% level of significance.

Further Exploratory Validations

As mentioned earlier, the computer had been programmed to play according to plans that were best responses to some plan of the participant and due to the instructions, this was common knowledge available to each participant. Hence we may regard second-order players to be the ‘best’ players in terms of game-theoretically rational thinking. If so, then the corresponding strategies of these players should be nearly perfect. Intuitively, we will have the least number of players from the second-order group committing a mistake while answering the final question regarding their belief corresponding to the computer’s future move. This fact is validated in the following Table.

Table 3: Summary statistics of the wrong answers according to ToM classifications of the players

Player type in terms of order of theory of mind	Number of participants falling in that group	Number of persons who gave wrong answers at some stage	Percentage
Zero-order	5	5	100%
First-order	27	21	77%
Second-order	16	9	56%

Discussion and Conclusion

In general, game-theoretic considerations lead to formation of typologies of players, which in turn can be validated statistically (cf. Rubinstein, 2014). Another approach to analyze a game-theoretic experiment is to identify it with a suitable logical system that expresses the experiment, followed by the construction of a computational cognitive model (cf. Ghosh, Meijering & Verbrugge, 2014).

In this paper, we follow a different method. We analyze an experiment about participants' behavior in a turn-taking game without going into the specifications of a game-theoretic model. Our aim is to develop robust domain-specific typologies of players. First, we classified the players by the probabilistic method of Latent Class Analysis, which is robust by construction. Furthermore, we used statistical techniques to validate the intuition behind those typologies. Secondly, we classified the players according to the order of theory of mind (ToM) they displayed; again, no game-theoretic considerations were taken into account. Once more, we validated the intuitions statistically.

What remains to be done is to check the interaction of the two typologies that were independently constructed. We find that 69% of the 'Second-order' players fall into the category of 'Expected and Learner' players. This validates our intuition that the set of 'Second-order' players and the set of 'Expected and Learner' players should not only have a non-empty intersection, but that the two sets should have quite an extensive section of players in common. We now aim to develop typologies that use the data on players' second decision times for validation purposes. One such typology would classify players into risk-taker versus risk-averse ones; another possible typology would distinguish competitive versus co-operative players. We also intend to design similar experiments to study various other possible typologies (for example, instinctive versus contemplative reasoners) and their interactions. The goal would be to build up a common perspective of typologies of players originating from various aspects of human strategic thinking.

Acknowledgments

We are grateful to Harmen de Weerd for his guidance with the statistical analysis. We thank the Indian Statistical Institute for awarding a travel grant to Tamoghna Halder for visiting the University of Groningen. Khyati Sharma's research stay in Groningen was supported by the Indo-European Research Training Network in Logic (IERTNiL) funded by the Institute of Mathematical Sciences, Chennai, the Institute for Logic, Language and Computation of the Universiteit van Amsterdam, and the Fakultät für Mathematik, Informatik und Naturwissenschaften, Universität Hamburg. Finally, we are grateful to the Netherlands Organization for Scientific Research for Vici grant NWO- 277-80-01, awarded to Rineke Verbrugge.

References

- Bateman, C., Lowenhaupt, R., & Nacke, L. E. (2011). Player typology in theory and practice. In: *Proceedings of DiGRA 2011 conference: Think design play*. Utrecht School of the Arts, Utrecht.
- Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, Princeton, NJ.
- Ghosh, S., Meijering, B., & Verbrugge, R. (2014). Strategic reasoning: Building cognitive models from logical formulas. *Journal of Logic, Language and Information*, 23(1), 1-29.
- Ghosh, S., Heifetz, A., & Verbrugge, R. (2015). Do players reason by forward induction in dynamic perfect information games? In: R. Ramanujam (ed.), *Proceedings of the 15th conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2015.
- Goodman, L. A. (1974). The analysis of systems of qualitative variables when some of the variables are unobservable. Part IA: modified latent structure approach. *American Journal of Sociology*, 79(5), 1179-1259.
- Hedden, T., & Zhang, J. (2002). What do you think I think you think? Strategic reasoning in matrix games. *Cognition*, 85(1), 1-36.
- Meijering, B., van Rijn, H., Taatgen, N. A., & Verbrugge, R. (2012). What eye movements can tell about theory of mind in a strategic game. *PLoS ONE*, 7(9), e45961.
- Meijering, B., Taatgen, N. A., van Rijn, H., & Verbrugge, R. (2014). Modelling inference of mental states: As simple as possible, as complex as necessary. *Interaction Studies*, 15(3), 455-477.
- Osborne, M. J. & Rubinstein, A. (1994). *A course in game theory*. The MIT Press, Cambridge, MA.
- Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52, 1029-1050.
- Raijmakers, M. E., Mandell, D. J., van Es, S. E., & Counihan, M. (2014). Children's strategy use when playing strategic games. *Synthese*, 191(3), 355-370.
- Rosenthal, R. (1981). Games of perfect information, predatory pricing and the chain store. *Journal of Economic theory*, 25(1), 92-100.
- Rubinstein, A. (2014). *A typology of players: Between instinctive and contemplative*. Tel Aviv, technical report.