

Using intrinsic complexity of turn-taking games to predict participants' reaction times

Jakub Szymanik (jakub.szymanik@gmail.com)

Institute for Logic, Language and Computation, University of Amsterdam

Ben Meijering (b.meijering@rug.nl)

Rineke Verbrugge (L.C.Verbrugge@rug.nl)

Institute of Artificial Intelligence, University of Groningen

Abstract

We study structural properties of a turn-based game called the Marble Drop Game, which is an experimental paradigm designed to investigate higher-order social reasoning. We show that the cognitive complexity of game trials, measured with respect to reaction time, can be predicted by looking at the structural properties of the game instances. In order to do this, we define complexity measures of finite dynamic two-player games based on the number of alternations between the game players and on the pay-off structure. Our predictions of reaction times and reasoning strategies, based on the theoretical analysis of complexity of Marble Drop game instances, are compared to subjects' actual reaction times. This research illustrates how formal methods of logic and computer science can be used to identify the inherent complexity of cognitive tasks. Such analyses can be located between Marr's computational and algorithmic levels.

Keywords: cognitive difficulty; strategic games; higher-order social reasoning; theory of mind

Introduction

In recent years, questions have been raised about the applicability of logic and computer science to model cognitive phenomena (see, e.g., Frixione, 2001; Stenning and Van Lambalgen, 2008; Van Rooij, 2008). One of the trends has been to apply formal methods to study the complexity of cognitive tasks in various domains, for instance: syllogistic reasoning (Geurts, 2003), problem solving (Gierasimczuk et al., 2012), and natural language semantics (Szymanik and Zajenkowski, 2010). It has been argued that with respect to its explanatory power, such analysis can be located between Marr's (1983) computational and algorithmic levels.

More recently, there has also been a trend to focus on similar questions regarding social cognition, more specifically, theory of mind. Especially, higher-order reasoning of the form 'I believe that Ann knows that Peter thinks ...' became an attractive topic for logical analysis (Verbrugge, 2009). Here, the logical investigations often go hand in hand with game theory (see, e.g., Osborne and Rubinstein, 1994). In this context, one of the common topics among researchers in logic and game theory has been backward induction (BI), the process of reasoning backwards, from the end of the game, to determine a sequence of optimal actions (Van Benthem, 2002). Backward induction can be understood as an inductive algorithm defined on a game tree. The BI algorithm tells us which sequence of actions will be chosen by agents that want to maximize their own payoffs, assuming common knowl-

edge of rationality. In game-theoretical terms, backward induction is a common method for determining sub-game perfect equilibria in the case of finite extensive-form games.¹

Games have been extensively used to design experimental paradigms aiming at studying social cognition (Camerer, 2003), recently with a particular focus on higher-order social cognition: the matrix game (Hedden and Zhang, 2002), the race game (Gneezy et al., 2010; Hawes et al., 2012), the road game (Flobbe et al., 2008; Raijmakers et al., 2013), and the Marble Drop Game (henceforth, MDG) (Meijering et al., 2010, 2011, 2012). All the mentioned paradigms are actually game-theoretically equivalent. They are all finite extensive-form games that can be solved by applying BI. As an example in this paper we will consider MDG (see Fig. 1).

Many studies have indicated that application of higher-order social reasoning among adults is far from optimal (see, e.g., Hedden and Zhang, 2002; Verbrugge and Mol, 2008). However, Meijering et al. (2010, 2011) report on a near ceiling performance of subjects when their reasoning processes are facilitated by, for example, a step-wise training. Still, an eye-tracking study of the subjects solving the game suggests that backward induction is not necessarily the only strategy used (Meijering et al., 2012).

We still do not know exactly what reasoning strategies² the subjects are applying when playing this kind of dynamic extensive form games. One way to use formal methods to study this question has been proposed by (Ghosh et al., 2010; Ghosh and Meijering, 2011): to formulate all reasoning strategies in a logical language, and compare ACT-R models based on each reasoning strategy with a subject's actual performance in a sequence of games, based on reaction times, accuracy and eye-tracking data. This corresponds to a study between the computational and algorithmic levels of Marr's (1983) hierarchy.

¹Backward induction is a generalization of the minimax algorithm for extensive form games; the subgame-perfect equilibrium is a refinement of the Nash equilibrium, introduced to exclude equilibria with implausible threats (Osborne and Rubinstein, 1994).

²The term 'strategy' is used here more broadly than in game theory, where it is just a partial function from the set of histories (sequences of events) at each stage of the game to the set of actions of the player when it is supposed to make a move. We are interested in human reasoning strategies that can be used to solve the cognitive problems posed by the game.




Figure 1: Examples of a zero-, first-, and second-order Marble Drop game. The blue marbles, on the left-hand side in the bins, are the participant’s payoffs and the orange marbles, on the right-hand side, are the computer’s payoffs. The marbles can be ranked from the lightest to the darkest. For each player, the goal is to get the white marble to drop into the bin with the darkest possible marble of their color. The participant controls the blue trapdoors (i.e., blue diagonal lines) and the computer controls the orange ones (the second set of trapdoors from the left). The participants are told that the computer aims at maximizing its pay-off. The dashed lines represent the trapdoors that both players should remove to attain the darkest possible marble of their color. See http://www.ai.rug.nl/~meijering/marble_drop.html for an interactive demo. Backward induction reasoning proceeds from the last decision, which in 1c is Player 1’s decision between the blue marbles in payoff-pairs C and D. Player 1 would decide to remove the left trapdoor because C contains the darker blue marble. Backward induction would then proceed with the second-to-last decision, which is Player 2’s decision between the orange marbles in payoff-pairs B and C. Player 2 would decide to remove the left orange trapdoor, because B contains the darker orange marble. Backward induction reasoning stops at the third-to-last decision, which is Player 1’s decision between the blue marbles in payoff-pairs A and B. Player 1 would remove the right blue trapdoor, because B contains the darker blue marble.

Here, we aim to tackle the problem from a somewhat more generic, complexity-theoretic viewpoint: we propose to study the problem on the computational level. Specifically, we will identify inherent, structural properties of the game that make certain MDG trials harder than others.

Alternation type

Every instance of a finite extensive form game can be presented as a decision tree. The second-order trials of MDG have the abstract tree form presented in Fig. 2.

How to approximate the complexity of a single instance of MDG? In the worst-case scenario, the backward induction algorithm, based on breadth-first search from the leaves of the tree upwards, will have to travel through all the nodes of the decision tree. Thus, it will find the rational solution (Nash Equilibrium) in time and space proportional to the number of nodes plus the number of edges in the tree, $O(|V| + |E|)$. However, the size of the tree does not seem to be a psychologically plausible complexity measure. To see this, consider two trees of equal size, but in the first one all the nodes are




Figure 2: Nodes s and u are controlled by Player 1. t is controlled by Player 2. If a player controls a node then in that node he can choose whether to go left, l , or right, r . Every leaf is labeled with the pay-offs for Players 1 and 2.

controlled by Player 1 while in the second tree, the players alternate. Obviously, the problem posed by the second tree is much more complex. This suggests that one of the key aspects of the problem is the structure of the move alternation in the game tree. Let us then categorize game trees with respect to such alternations. In the following, we restrict the analysis to two-player games, although it would be possible to extend the ideas to finite dynamic games for more than two players.

Definition 1 *Let us assume that the players $\{1, 2\}$ strictly alternate in the game; Let player $i \in \{1, 2\}$. Then:*

- In a Λ_1^i tree, all the nodes are controlled by Player i .
- A Λ_{k+1}^i tree, a tree of k -alternations for some $k \geq 0$, starts with a Player i node.³

For instance, the tree in Fig. 2 is Λ_3^1 , a 1-game tree of 2 alternations, because Player 1 has the first move at the root, followed by an alternation of Player 1 to Player 2 and another alternation of Player 2 to Player 1.

Pay-off structure and cognitive difficulty

From the psychological perspective, it seems really crucial to take pay-offs into account when comparing the difficulty of particular MDG tasks. For instance, the two trees from Fig. 3 are Λ_3^1 , because they both start with Player 1 and both have two alternations, from Player 1 to Player 2 and back again. However, clearly, the first game, represented by T_1 , is much easier for Player 1 than the second game, represented by T_2 . In the first game it is enough for Player 1 to realize that 999 is the highest possible pay-off, and then he can instantly move left and finish the game.

To explain the eye-tracking data of the subjects solving the Marble Drop game, Meijering et al. (2012) suggest that subjects may be using forward reasoning with backtracking (henceforth FRB), based on statistical analysis of eye gaze sequences. For instance, in the game from Fig. 1c, Player 1 will find out that B contains the darkest blue marble. He has to ask himself whether that marble is attainable. In other words, he

³From the computational complexity theory perspective, this corresponds to a hierarchy of computational problems of increasing complexity (see, e.g., Arora and Barak, 2009).




Figure 3: Two Λ_3^1 trees.

has to reason about whether Player 2 would remove the left orange trapdoor. Therefore, Player 1 has to look at the orange marbles in bins B, C and D to find out that bin D contains Player 2's darkest orange marble. The reasoning continues with Player 1 asking himself whether Player 2 thinks that her orange marble in bin D is attainable. In other words, Player 1 has to reason about whether she thinks that he would remove the right blue trapdoor of the rightmost set of trapdoors. Player 1 knows that he would not remove that trapdoor, but that he would remove the left one instead. He also knows that she is aware of this, as both players are aware of each other's goals. Therefore, Player 1 knows that Player 2 knows that her darkest orange marble in D is unattainable. Therefore, Player 1 has to go back to the second decision point (i.e., the orange trapdoors). There, Player 2 would compare the orange marbles in B and C and decide to remove the left orange trapdoor, because the orange marble in B is the darkest orange marble that she can still attain. To conclude, Player 1 knows that his darkest blue marble in B is attainable, and will thus remove the right blue trapdoor of the leftmost set of trapdoors.

As it is relatively hard to conclude from the eye-tracking data whether subjects apply exactly the above described forward reasoning with backtracking, we propose an orthogonal idea. We aim to identify the properties of the games that make certain trials harder than others and see whether such an explanation is congruent with forward reasoning plus backtracking. In order to do that, we put forward the following definitions. The idea here is that subjects may be looking for the highest possible pay-off and then try to reach it.

Definition 2 A game T is generic, if for each player, distinct end nodes have different pay-offs.

Note, for instance, that the game in Figure 1c is generic: the four bins contain marbles of four different hues of blue

and four different hues of orange.

Definition 3 Suppose $i \in \{1, 2\}$. If T is a generic game tree with the root node controlled by Player i and n is the highest possible pay-off for Player i , then T^- is the minimal subtree of T containing the root node and the node with pay-off n for Player i .

To illustrate this definition, Figure 4 shows the restricted T^- trees for the two trees shown in Figure 3.

Hypothesis 1 Let us take two MDG trials T_1 and T_2 . T_1 is easier for participants than T_2 if and only if T_1^- is lower in the tree alternation hierarchy than T_2^- .

Hypothesis 1 takes into account pay-off structures. According to it, the first tree from Fig. 3, T_1 , should be easier for participants than the right tree, T_2 , as T_1^- is a Λ_1^1 tree while T_2^- is still Λ_3^1 , see Fig. 4. Moreover, it is possible that some subjects may try to apply the procedure iteratively: check if the maximum pay-off is reachable, if not then check for the second-best pay-off, and so on.




Figure 4: The maximum pay-off restricted trees corresponding to the trees in Fig. 3.

As an additional question, we ask whether the following predictions agree with the proposal of Meijering and colleagues (Meijering et al., 2012) that the subjects in the game are applying forward reasoning, with backtracking when necessary (FRB). First of all, why would subjects ever apply FRB?

Hypothesis 2 For an average random game with 3 decision points structured as the Λ_3^1 game of Figure 2, the forward reasoning plus backtracking algorithm needs fewer computation steps to yield a correct solution than backward induction.

Furthermore, if subjects used forward reasoning, then we could observe the following by running FRB algorithm on the game trees:

Hypothesis 3 Let us take two MDG trials T_1 and T_2 . The forward induction with backtracking algorithm yields a correct solution for T_1 faster than for T_2 if and only if T_1^- is lower in the tree alternation hierarchy than T_2^- .

Experimental results

To experimentally corroborate our hypotheses, we analyzed performance and reaction time data from (Meijering et al.,

2012). Twenty-three first-year psychology students (14 female) with a mean age of 20.8 years (ranging from 18 to 24 years) participated in the experiment and were asked to solve Marble Drop trials, in the sense that they had to make a decision ‘left’ or ‘right’ at the first decision point. All experimental game trials had payoff structures that required Player 1 to reason about the decision at each of the three decision points, structured as the Λ_3^1 game of Figure 2. Therefore, the experiment was constructed in a way to be diagnostic for second-order theory of mind (see Meijering et al., 2012, for more information on the experimental design).

We divided experimental trials into two sets: **Accessible** ones, in which the highest possible pay-off for Player 1 is obtainable for him and **Inaccessible** ones, where his highest possible pay-off is not obtainable. For example, the game of Figure 1c is accessible, because Player 1 can reach the marble of the darkest hue of blue, which is located in bin B, by opening the right trapdoor; after all, Player 2 will also choose to stay there. Note that in general, if T_1 represents an accessible game and T_2 an inaccessible one, then T_1^- is lower in the alternation hierarchy than T_2^- .

Therefore, according to Hypothesis 1, our prediction was that the shortest reasoning times will be recorded in the condition “Accessible”, where the highest pay-off was obtainable for Player 1.

Furthermore, by simulating forward reasoning with backtracking on experimental trials and computing the number of reasoning steps, we investigated hypotheses 2 and 3. Again, our prediction was that the number of steps should be smaller in “accessible” cases, where the highest-possible pay-off for Player 1 was obtainable.

Hypothesis 1: pay-offs and alternation type

To investigate the first hypothesis, we compared reaction times (RTs) in games in which the highest payoff was accessible against RTs in games in which the highest payoff was not accessible. The RTs were log-transformed to approximate the normal distribution.

A paired-samples t-test indicated a significant (within-subjects) difference, $t(12) = 4.07, p < .01$. The RTs decrease if the maximum payoff is accessible, which can be seen in Figure 5.

Hypothesis 2: simulating the algorithms

When looking at all possible payoff-structures in Marble Drop games with two alternations (or three decision points), we implemented the forward reasoning plus backtracking algorithm as a set of heuristics based on several cases that can occur in the Marble Drop game; we used the same algorithm that we derived in (Meijering et al., 2012) from the participants’ eye-tracking data.⁴

When using the algorithm on all 576 possible pay-off structures, we see that forward reasoning with backtracking in

⁴Thus, we did not use a generic implementation of forward reasoning with backtracking that would work for any possible game tree.




Figure 5: Players’ reaction times with respect to accessibility, namely the attainability of the highest payoff for Player 1.

general requires fewer steps than backward induction, e.g., in 288 cases only 1 step is enough. More specifically, forward reasoning with backtracking requires on average 3 steps, whereas backward induction would always require 6 steps, irrespective of payoff structure. Table 1 provides a cross-table of payoff structures and number of steps. This simulation supports our Hypothesis 2.

Table 1: Cross-table of payoff structures and the necessary number of steps when using forward reasoning with backtracking.

# of steps	1	2	4	5	6	8
# of payoff structures	288	72	48	56	16	96

These simulation results imply that, on average, it pays off to use a forward reasoning strategy. In fact, Meijering et al. (2012) found a strong prevalence of forward reasoning with backtracking, even though participants were presented with a subset of hard-to-solve games in which backward induction would actually be more efficient on average. However, participants did not know that they were presented with this particular subset of very difficult games.

Hypothesis 3: FRB and structural complexity

The implementation of the forward reasoning plus backtracking (FRB) algorithm was applied to the subset of actually presented experimental games to determine the number of reasoning steps required for each game. In the following analyses, number of steps was included as a predictor of the reaction times. We label the factor simply as ‘forward reasoning with backtracking’.

The log-RTs were analysed by means of linear mixed-effects (LME) models (Baayen et al., 2008) to account for random effects of participants and unequal numbers of observations across all experimental conditions. Traditional (repeated measures) ANOVAs could not be performed as they

require equal numbers of observations.

Fitting LMEs on the log-transformed reaction times, we see that forward reasoning plus backtracking (FRB) is a good predictor. The model with FRB cannot be rejected in favor of a simpler model without FRB as a predictor, $\chi^2(1) = 8.4$, $p = 0.004$. We discuss the best model below.

Again, the reaction times significantly decrease if the maximum Player 1 payoff is accessible (Table 2a). In case of games in which the maximum payoff is not accessible, the reaction times do not significantly increase with each additional reasoning step (Table 2b). Those games require in between 6 and 8 reasoning steps, which is too small a difference to find a significant effect on the RTs. In contrast, the RTs do significantly increase with each additional reasoning step in games in which the maximum payoff is accessible (Table 2c).

Table 2: Output of full-factorial linear mixed-effects model with factors Accessibility (A), Steps of forward reasoning with backtracking (FRB).

Parameter	Estimate	St. Error	t-value	p-value
a) Accessible	-0.689147	0.271256	-2.54	.000
b) FRB	0.008767	0.034930	0.25	.418
c) A:FRB	0.084336	0.037277	2.26	.000

Discussion

We have investigated the structural properties of the Marble Drop Game, an experimental paradigm designed to study higher-order social reasoning. Using theoretical approaches from logic and complexity theory, we identified inherent properties of the game trials responsible for the cognitive difficulty of the task. Meijering and colleagues' (2012) reaction time data can be explained by looking at the alternation type and pay-off distribution of the particular game items. It turned out that the game items are harder if the maximum possible pay-off for Player 1 is not accessible for him. This observation is consistent with the assumption that participants were mostly applying forward reasoning with backtracking to solve the games. By simulating forward reasoning with backtracking on the experimental items, we have shown that the reaction times and the number of necessary comparisons significantly decrease if the maximum Player 1 payoff is accessible. As MDG is game-theoretically equivalent to many other experimental paradigms making use of turn-based games (see, e.g., Hedden and Zhang, 2002; Gneezy et al., 2010; Hawes et al., 2012; Flobbe et al., 2008; Raijmakers et al., 2013), we would expect that our results generalize to those cases.

One could wonder why the subjects did not use backward induction in the first place, as it is the method that always delivers the optimal pay-off (Osborne and Rubinstein, 1994). One possible answer is that they avoided backward induction in order to simplify the underlying reasoning. Recall, that

while backward induction reasoning always takes 6 steps in the Marble Drop game with 3 decision points, forward reasoning and backtracking takes on average only 3 steps, corresponding with the phenomenon that T^- is usually lower in the tree alternation hierarchy than T itself. Moreover, iterating the forward reasoning strategy by backtracking in case the highest pay-off is not obtainable will finally lead to the optimal solution. Therefore, some subjects may choose to use that strategy to avoid higher-order reasoning, even though keeping the intermediate results in mind during backtracking is expected to tax working memory more than applying backward induction.

Subjects may as well use other heuristics that do not guarantee reaching the prescribed backward induction result, namely a Nash equilibrium of the game. For instance, as suggested by Hedden and Zhang (2002), subjects may assume that their opponents are playing according to some fixed patterns. Instead of assuming that the opponent is rational and correctly predicts Player 1's choice at the last decision point, Player 1 may take his opponent to be risk-averse or risk-taking. Such heuristics, essentially based on considering sub-trees of the initial game-tree, will also lead to simplified reasoning.

Of course, assuming that the opponent is of some specific type changes the game drastically and can lead to a very bad outcome, in case of wrong judgement of the other player's type. Still, people notoriously apply similar heuristics in strategic situations, for example, when joining a poker table, many players try to evaluate whether the opponents play 'loose' or 'tight'.⁵ An important question is what are the good alternative strategies. They should be not only easy to compute for people but also relatively safe to apply. It seems that the forward reasoning plus backtracking strategy in MDG might be a cognitively attractive strategy for people asked to solve turn-based games. First of all, it does not ask for the second-order social reasoning that is known to be very hard even for many adults (Verbrugge, 2009), and moreover, on average it demands fewer comparisons. One may even think that competent players know a collection of various strategies and their strategic abilities could be partially equated with the skill of choosing the right one, i.e., a strategy that may be safely applied in a given context to simplify the underlying reasoning.

Outlook

Inspired by the logical study of backward induction and the cognitive science experiments with the Marble Drop Game, we investigated structural properties of turn-taking dynamic games and we provided a more refined analysis of the complexity of particular game trials, which takes into account alternation type of the game and pay-off distribution. We com-

⁵A similar phenomenon is well-recognized in natural language semantics. People often shift the meaning of sentence ϕ from $\llbracket\phi\rrbracket$ to a more restricted meaning $\llbracket\psi\rrbracket \subseteq \llbracket\phi\rrbracket$. And again, one of the factors triggering such meaning-shifts might be related to the computational complexity of ϕ (see, e.g., Szymanik, 2010).

pared our predictions to actual reaction time data from (Meijering et al., 2012).

Of course, there are many further topics to be resolved. For instance, it would be interesting to extend our analysis to account for imperfect information games. Also it would be fruitful to explore connections with various related logical formalisms and to investigate further epistemic phenomena. In parallel, we would like to confront Hypotheses 1 and 3 with the available eye-tracking data from (Meijering et al., 2012), as well as with eye-tracking data to be gathered from a wider class of turn-based two-player games. Moreover, we plan to investigate other reasonable reasoning strategies that subjects may successfully adapt in game-plays.

Acknowledgments

The authors are grateful for the support of Vici Grant NWO-277-80-001 awarded to RV. JS also acknowledges NWO Veni Grant 639.021.232.

References

- Arora, S. and Barak, B. (2009). *Computational Complexity: A Modern Approach*. Cambridge University Press, New York, NY, USA, 1st edition.
- Baayen, R., Davidson, D., and Bates, D. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59:390–412.
- Van Benthem, J. (2002). Extensive games as process models. *Journal of Logic, Language and Information*, 11(3):289–313.
- Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, New Jersey.
- Flobbe, L., Verbrugge, R., Hendriks, P., and Krämer, I. (2008). Children’s application of theory of mind in reasoning and language. *Journal of Logic, Language and Information*, 17(4):417–442.
- Frixione, M. (2001). Tractable competence. *Minds and Machines*, 11(3):379–397.
- Geurts, B. (2003). Reasoning with quantifiers. *Cognition*, 86(3):223–251.
- Ghosh, S. and Meijering, B. (2011). On combining cognitive and formal modeling: A case study involving strategic reasoning. In Van Eijck, J. and Verbrugge, R., editors, *Proceedings of the Workshop on Reasoning About Other Minds: Logical and Cognitive Perspectives (RAOM-2011)*, Groningen. CEUR Workshop Proceedings.
- Ghosh, S., Meijering, B., and Verbrugge, R. (2010). Logic meets cognition: Empirical reasoning in games. In Boissier, O., Fallah-Seghrouchni, A. E., Hassas, S., and Maudet, N., editors, *MALLOW*, volume 627 of *CEUR Workshop Proceedings*, Lyon. CEUR-WS.org.
- Gierasimczuk, N., van der Maas, H., and Raijmakers, M. (2012). Logical and psychological analysis of deductive Mastermind. In Szymanik, J. and Verbrugge, R., editors, *Proceedings of the Logic & Cognition Workshop at ESSLLI 2012, Opole, Poland, 13-17 August, 2012*, volume 883 of *CEUR Workshop Proceedings*, pages 1–13, Opole. CEUR-WS.org.
- Gneezy, U., Rustichini, A., and Vostroknutov, A. (2010). Experience and insight in the race game. *Journal of Economic Behavior and Organization*, 75(2):144 – 155.
- Hawes, D. R., Vostroknutov, A., and Rustichini, A. (2012). Experience and abstract reasoning in learning backward induction. *Frontiers in Neuroscience*, 6(23).
- Hedden, T. and Zhang, J. (2002). What do you think I think you think?: Strategic reasoning in matrix games. *Cognition*, 85(1):1 – 36.
- Marr, D. (1983). *Vision: A Computational Investigation into the Human Representation and Processing Visual Information*. W.H. Freeman, San Francisco.
- Meijering, B., Van Maanen, L., Van Rijn, H., and Verbrugge, R. (2010). The facilitative effect of context on second-order social reasoning. In Catrambone, R. and Ohlsson, S., editors, *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, pages 1423–1428, Austin (TX). Cognitive Science Society.
- Meijering, B., van Rijn, H., Taatgen, N. A., and Verbrugge, R. (2012). What eye movements can tell about theory of mind in a strategic game. *PLoS ONE*, 7(9):e45961.
- Meijering, B., Van Rijn, H., and Verbrugge, R. (2011). I do know what you think I think: Second-order theory of mind in strategic games is not that difficult. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*, pages 2486–2491, Boston. Cognitive Science Society.
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. The MIT Press, Cambridge, MA.
- Raijmakers, M. E., Mandell, D. J., Es, S. E., and Counihan, M. (2013). Children’s strategy use when playing strategic games. *Synthese*.
- Van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science: A Multidisciplinary Journal*, 32(6):939–984.
- Stenning, K. and Van Lambalgen, M. (2008). *Human Reasoning and Cognitive Science*. The MIT Press, Cambridge, MA.
- Szymanik, J. (2010). Computational complexity of polyadic lifts of generalized quantifiers in natural language. *Linguistics and Philosophy*, 33:215–250.
- Szymanik, J. and Zająkowski, M. (2010). Comprehension of simple quantifiers. Empirical evaluation of a computational model. *Cognitive Science: A Multidisciplinary Journal*, 34(3):521–532.
- Verbrugge, R. (2009). Logic and social cognition: The facts matter, and so do computational models. *Journal of Philosophical Logic*, 38(6):649–680.
- Verbrugge, R. and Mol, L. (2008). Learning to apply theory of mind. *Journal of Logic, Language and Information*, 17(4):489–511.