Trust as Theory of Mind

Bachelor's Project Thesis

Hanna Marleen Havinga, s2657139, H.M.Havinga@student.rug.nl,
Supervisors: Dr H.A. de Weerd & Prof Dr L.C. Verbrugge

Abstract: When humans play the prisoner's dilemma it is observed that sometimes they cooperate after a defect-defect situation has occurred. This paper investigated whether this can be explained by the ability of people to reason about the beliefs and goals of others, also known as theory of mind. This is tested through a simulation with agents who use different orders of theory of mind in which the level of cooperation is measured. The results show that the there is a significant difference between the games in which agents use different orders theory of mind. The data suggests that when a higher order theory of mind is used there is less cooperative behaviour. Therefore we conclude that theory of mind does not cause cooperation by itself.

1 Introduction

1.1 Theory of mind

Theory of mind (ToM) (Premack and Woodruff, 1978) is the ability to attribute mental states to other people. This means that someone is able to understand that others have beliefs about the world that may be different from their own. When we speak about ToM, we often identify different orders of ToM. Zero-order ToM means that the person only reasons from their own point of view; what they believe to be true about the world. An example of this reasoning is, "I think the ball is in the basket". First-order ToM is the ability to reason about what someone else may believe about the world, like "I think that they think that the ball is in the basket". With second-order ToM there comes the understanding that someone else also has theory of mind and that another person can, for example, also attribute mental states to you (such as: "I think that they think that I think the ball is in the basket"). This line of thinking can be repeated infinitely and that is referred to as higher-orders theory of mind.

Most adults use theory of mind up to the fourth order (Kinderman, Dunbar, and Bentall, 1998). However, in many strategic game settings, like the prisoner's dilemma, people are found to rely on low orders of ToM and to be slow to increase their order of ToM reasoning (Goodie, Doshi, and Young, 2012; Camerer, Ho, and Chong, 2004). Theory of mind has been suggested to be useful in cooperation (Tomasello, 2009; Tomasello, Carpenter, Call, Behne, and Moll, 2005), we discuss this further in section 1.3.

	B cooperates	B defects
A cooperates	Both 3	B = 5, A = 0
A defects	A = 5, B = 0	Both 2

Table 1.1: Payoff matrix of the prisoner's dilemma. Player A decides the row, while player B decides the column.

1.2 Prisoner's dilemma

The prisoner's dilemma is a situation used in game theory to look at individual and group gain. In this situation, two agents have the ability to betray or cooperate with the other. When they both betray, they achieve a small reward. When they both cooperate, they achieve a moderate reward and when one betrays and the other cooperates, it results in the one betraying to achieve the largest reward and the one being betrayed to get the smallest reward (Kuhn, 2017).

In Table 1.1 an instance of the prisoner's dilemma is portrayed. Here we have two agents who play the game, agent A and agent B, who both have two moves they can make: to cooperate (C) or to defect (D). When they both cooperate, they both gain 3 points and when they both defect, they gain 2 points. When one agent cooperates and is betrayed, it gains no points and the one who betrays gains 5 points.

When agent B is going to defect, it is better for agent A to defect than to cooperate, since the former leads to a gain of 2 while the latter results in no gain. If agent B were to cooperate it would again be better for agent A to defect, since this leads to a gain of 5 while cooperating leads to a gain of only 3. The same reasoning goes for agent B. However, if both agents cooperate the reward is higher for both of them than when they both defect.

This leads to the dilemma of achieving higher rewards with a higher risk of being betrayed, or achieving lower rewards with no risk.

1.3 Previous research

When people play the iterated version of the prisoner's dilemma they usually fall into a situation in which both people defect (D-D) and keep defecting. It has been observed that, after a stable D-D situation, people sometimes make a cooperative move (Monterosso, Ainslie, Mullen, and Gault, 2002). This is surprising because when a person cooperates and their opponent defects they will gain the lowest possible reward. This means that this is the worst move that can be made when we look at only one game. However, if on the long term making a cooperative move will get the opponent to also cooperate, this will lead to higher rewards for both players, since the rewards gained from both agents cooperating is higher than from both defecting (see Table 1.1).

The goal of both agents in the prisoner's dilemma is gaining the highest rewards possible and by cooperating both agents can get a higher reward than if they were to defect. This makes the prisoner's dilemma a good situation for shared intentionality (Tomasello, 2009; Tomasello et al., 2005). By using higher-order ToM shared intentionality can occur, where both agents assume that the other will work with them towards the common goal of achieving higher rewards for everyone. This way higher-order ToM would lead to more cooperation.

1.4 Research question and hypothesis

In this paper we focus on the following question: is it theory of mind that leads to cooperation? And more specifically, can higher orders of theory of mind get the agent from a D-D situation to a C-C situation?
We expect theory of mind to lead to more cooperation. This could result in more cooperation with every additional level theory of mind or a sudden rise in cooperation at a certain level theory of mind.

In Section 2 of this paper we will take a look at the environment that was created and the experiment that was run in it. We will discuss the results from this experiment in Section 3 and in Section 4 we will look at the discussion, the conclusion and possible future research.

2 Method

To run our experiment we have built an environment in which agents play an iterated version of the prisoner's dilemma. In this section, we will explain the environment that was built, the model that the agents are built on and the algorithm that was used.

2.1 Environment

2.1.1 Agent Simulation

For this experiment we use agent-based simulation. This is a simulation technique in which agents, entities that act autonomously, interact with each other. When several agents interact this can lead to complex group behaviour. A reason we use agent-modeling is because it is easy to change what mental abilities an agent has. This way we can test how agents who are capable of using different orders ToM interact with each other. Another reason to use agents instead of human subjects is that we can isolate ToM from other mental abilities. This has allowed us to be sure that the behaviour we observe is a result of the level of ToM that an agent uses.

	B cooperates	B defects
A cooperates	Both 100	B = 101, A = 1
A defects	A = 101, B = 1	Both 2

Table 2.1: Adjusted payoff matrix of the prisoner's dilemma used in the simulation experiment.

2.1.2 Simulation-Theory of Mind

The ToM that our agents use is based on simulation-theory of mind. Simulation-theory of mind says that people use theory of mind by taking the perspective of someone else, while assuming that the other person thinks like you (Barlassina and Gordon, 2017). Therefore the way you think is used as a simulation for how someone else thinks, and you expect the other to react the same as how you would react if you were to be in their situation.

2.1.3 Agents playing the Prisoner's Dilemma

The payoff matrix that we use in our environment isn't a standard one. We have adjusted the values of the rewards to encourage cooperation, even though the fundamental structure of the prisoner's dilemma remains true, where an agent always receives a higher reward by defecting than it would by cooperating. The way in which cooperation is encouraged by this matrix is that the difference between the rewards to cooperate and defect are small and the difference between a C-C situation and a D-D situation are very large, as can be seen in Table 2.1.
For example, if agent B were to cooperate, there would only be 1 point difference in the gain for agent A between cooperating or defecting, so the reward for betraying your opponent is very small. However, the difference between both agents cooperating and both agents defecting is 98 points. With these rewards it becomes very important to make sure that both agents cooperate at the same time, since a stable C-C situation leads to very high rewards in comparison to a stable D-D situation.

2.2 Explanation of the model

In the iterated version of the prisoner's dilemma the same agents consecutively play multiple rounds of the prisoner's dilemma, where one round consists of both agents making a move (cooperate or defect) as output. The game state is what has occurred in the previous round, which is the output of both agents in the last round.

In our setting, agents believe that their opponent may change its behaviour in response to the agent's actions. When an agent considers what move to make for the current round, it also considers how its actions influences several rounds into the future. When you think 3 steps into the future there are $2^3 = 8$ plans that you can follow. A plan would for example be [C,C,D]. For the first action in the plan (the one the agent would play in the current turn) it calculates what its gain would be when the opponent defects or cooperates. To that, the agent adds what it will gain in the next turns when following the actions of the plan, but the gain from every turn that's a step further into the future gets multiplied by a discount factor.
For each of these plans the agent calculates what the gain would be in the current and next three turns. In the current turn it plays the first action of the plan with the highest gain.

Every agent has the ability to use ToM to a certain order. The base of every agent is the zero-order theory of mind strategy, in which it only reasons from its own perspective. Here the probability of the opponent cooperating is calculated by taking the percentage of times the opponent has chosen to cooperate after the same game state has occurred. For example, if the current game state is [D-D] and our opponent has played 'cooperate' 25% of the time following a [D-D] situation, then the zero-order agent will believe that there is a 25% chance that the opponent will play 'cooperate' now.

For the first-order ToM strategy the agent simulates a zero-order agent, in which the simulated agent receives the same parameters and knowledge as the agent, but the memory of the game states is given from the perspective of the opponent. In this way an opponent with zero-order ToM is simulated. The agent observes how the simulated agent would respond in this situation. It then integrates this knowledge into the probability of the opponent cooperating, that was calculated with the zero-order strategy, based on its confidence in the first-order strategy. For the second- and third-order strategies the agents use the same method as the first-order strategy, by simulating an agent that has first-order and second-order theory of mind respectively. When a strategy predicts the move of the opponent correctly the confidence that the agent has in that strategy increases.

2.3 Algorithm in words

Every agent is initialized with several parameters; the number of prospective thinking steps (S), discount factor (δ), learning speed (λ), the memories of the agent's actions and the opponent's actions (these can be empty), and the confidence levels that the agent has for every level of ToM that it is capable of using.

The number of prospective thinking steps determines how many rounds in the future the agent considers when deciding on the best course of action. It can influence the agent's behaviour since the move that is made in this round will affect how the opponent reacts in the next round.

The discount factor shows how patient the agent is, how willing to lose this round to win the next. When the discount factor is zero, the agent only cares about the gain of the current round and no future rounds' possible gains are taken into consideration. When the discount factor is one this means that the gains of all future rounds are just as important as the gain of the current round. When the discount factor lies somewhere between zero and one, the importance of a round's gain gradually decreases with each step into the future.

For every strategy, except zero-order, that an agent is capable of using, it has a confidence-level. The confidence level shows how good this strategy is at predicting the opponent's next move. The confidence increases when the opponent's move has been predicted correctly and decreases when it has been predicted incorrectly. The reason this isn't true for the zero-order strategy is because it doesn't simulate an opponent, but merely observes previous game states. The predictions of the other strategies are integrated into the zero-order predictions, based on how much confidence the agent has in the strategy.

The learning speed shows how much influence the last round has on the agents confidence in its strategies. For example, when the learning speed is high and the first-order strategy has predicted the opponent's next move incorrectly, the confidence the agent has in its first-order strategy will decrease a lot. How this influence works exactly can be seen in Algorithm 2.1.

Algorithm 2.1 Update Confidence

Require: Predicted move
Ensure: Confidence updated based on performance
  if predicted move == opponent move **then**
    $C \leftarrow \lambda + (1-\lambda) * C$
  else
    $C \leftarrow (1-\lambda) * C$
  end if
  return

Since the number of prospective thinking steps is static, all possible plans are generated at initialization. When a round is started the agent will check the gain of every plan it could follow and choose the first move of the plan with the highest gain as its output.

Algorithm 2.2 Calculate Gain

Require: policy Π with actions A; past Game States
Ensure: result = gain of policy
  $P\_opp_C \leftarrow CalculateProbability()$
  $P\_opp_D \leftarrow 1 - P\_opp_C$
  if Last A in Π **then**
    result $\leftarrow P\_opp_C * gain(A, C) + P\_opp_D * gain(A, D)$
  else
    $GameStates_C \leftarrow$ add(GameStates, [A, C])
    $GameStates_D \leftarrow$ add(GameStates, [A, D])
    result $\leftarrow P\_opp_C * gain(A, C) + \delta$*CalculateGain($\Pi, GameStates_C$) $+ P\_opp_D * gain(A, D) + \delta$*CalculateGain($\Pi, GameStates_D$)
  end if
  return result

The gain of a plan is calculated as follows and can be seen in Algorithm 2.2. First, the probability of the opponent cooperating in the current game state is calculated. From this the probability of the opponent defecting is also calculated. Then the next action in the plan is taken. If this is not the last action in the plan the result is calculated as follows. In the next turn the opponent will either cooperate or defect. For both these possible scenarios the agent makes a hypothetical collection, which contains all the past game states and a new game state in which the opponent either cooperated or defected. For both scenarios the agent calculates what the gain would be if the opponent were to take one action and the agent were to play the current action in its plan. This gain is then multiplied by the chance that the opponent will actually take that action. Since this is not the last action in the plan, this function will be called again to look at the next action in the plan, but this time a hypothetical collection of game states will be sent along that contains the move the opponent could have taken. The output of the function that was recursively called is multiplied by the discount factor, since future gains are less important. This is calculated for both the opponent cooperating and defecting and the results of those two calculations are added together.

If the current action is the last action in the plan then the result is only the gain of the opponent making one move and the agent making the last move in its plan multiplied by the chance of the opponent actually making that move. This is again done for both the opponent cooperating and defecting.

The probability that an opponent will cooperate is calculated as follows and can be seen in Algorithm 2.3. First the zero-order strategy is calculated. Here it is checked whether the current game state has happened before in past rounds of the game. If it did occur this is recorded as well as whether after that game state the opponent cooperated. When the current game state has not occurred before, the probability of the opponent cooperating becomes 0.5. If the current game state did occur in the past the probability of the opponent cooperating is calculated by dividing the amount of cooperation by the amount of previous instances of the current game state. This indicates how often the opponent has cooperated in the past when faced with the same situation.

A zero-order agent only uses the above calculation. When an agent is capable of using higher orders of

Algorithm 2.3 Calculate Probability

Require: current Game State (cGameState); past Game States (GameStates)
Ensure: result = probability opponent will cooperate

```
for all GameStates in Memory do
  if cGameState == GameState then
    equal++
    if Opponent cooperates after GameState then
      cooperation++
    end if
  end if
end for
if No equal GameStates found then
  probability ← 0.5
else
  probability ← cooperation/equal
end if
for all x in orders ToM above Zero-order do
  SimAgent ← Simulate (x-1)-order agent with opponent's memory
  if SimAgent cooperates then
    probability ← (1-Confidence_x)*probability + Confidence_x
  else
    probability ← (1-Confidence_x)*probability
  end if
  Save move predicted by x-order ToM
end for
return probability
```

ToM, the following calculation is used. When we take first-order strategy as an example, the agent simulates an opponent with zero-order strategy by initializing a zero-order agent which will be given the same parameters as our agent, but with the game states from the perspective of the opponent. This simulated agent is then tasked to play the round and the output is observed by the playing agent. If the simulated opponent cooperates this is integrated into the probability that the actual opponent will cooperate based on the confidence that the agent has in the first-order strategy. If the simulated opponent defects this is also integrated by lowering the probability that the opponent will cooperate based on the confidence in this strategy.
This is done for every higher-order ToM, by simulating an opponent with a ToM that is one order lower than the agent's ToM. When all orders that the agent is capable of using are integrated into the probability of the opponent cooperating then the final calculated probability is used in the calculation of the gain of a policy.

2.4 Research protocol

An Iterated Prisoners dilemma game contains 120 rounds. The first 17 of these rounds are random, meaning that in each of these rounds the moves that the agents make are decided randomly. Next, each agent observes 3 rounds in which both agents defect. This is done to ensure that we can observe the behaviour that occurs after a stable D-D situation. Finally, the agents play 100 rounds freely in which we observe how often the agents cooperate and how often they cooperate at the same time.

All agents were given the same parameters in every game. The agents think 3 steps ahead, the discount factor they use is 0.2 and the learning speed is 0.6. Both agents start with zero confidence in all orders of ToM, which can change over the course of the game. The number of thinking steps was kept low to keep the simulation within reasonable time and the discount factor is low to make sure that the future rounds that the agent considers are still important. The agents have a memory of 20 rounds.

A one-way ANOVA will be used if the results are normally distributed. If the results are not normally distributed a non-parametric Kruskal-Wallis test will be used to determine whether there is a significant difference between groups.

3 Results

An agent pair consists of two agents of which the highest-order theory of mind they can use lies somewhere between zero- and third-order. Every agent pair plays 1000 games with 120 rounds each. A round consists of each agent making a move (cooperate or defect).

For every game the total amount of cooperative moves was recorded. The results have been summarized in Figure 3.1. In this figure, we show the number of times that an agent of the agent pair cooperated during the games that were played. Both agents' moves were recorded as separate cooperative moves, which means that when there are 100 rounds in a game the maximum number of cooperative moves is 200. Within this figure, the position of the agents are interchangeable, which means that, for example, the results of the 3-2 agent pair is the same as the 2-3 agent pair.

At first glance we see that the 0-0 agent pair has on average cooperated far more than the other agent pairs and that the other agent pairs have roughly the same amount of cooperation, apart from the 3-3, 3-2 and 2-2 agent pairs, which seem to have cooperated less on average.

The highest amount of cooperation within a game was 180 cooperative moves, which was produced by a 0-0 agent pair, while the lowest amount of cooperation was 0 cooperative moves, which was produced by the 3-3, 3-2 and 2-2 agent pairs.

Both the 0-0 and 1-1 agent pairs have clusters of outliers; when we look at the recorded data we have found that within these specific games the amount of C-C situations, in which both agent cooperate at the same time, make up (almost) all cooperative moves. For example, a few games of the 0-0 agent pair have in total 180 cooperative moves and 90 C-C situations, which means all cooperative moves were made at the same time as the opponent's. When taking a look at the specific rounds of a game played by a 0-0 agent-pair with 180 cooperative moves we have found that within such games the agents cooperate nine times in a row and then both defected for one turn before turning back to cooperation. The only thing that made these agents different from each other was their memory, which we expected to show some pattern that would explain this behaviour, like more cooperative moves or a mirrored memory in which both agents experi-

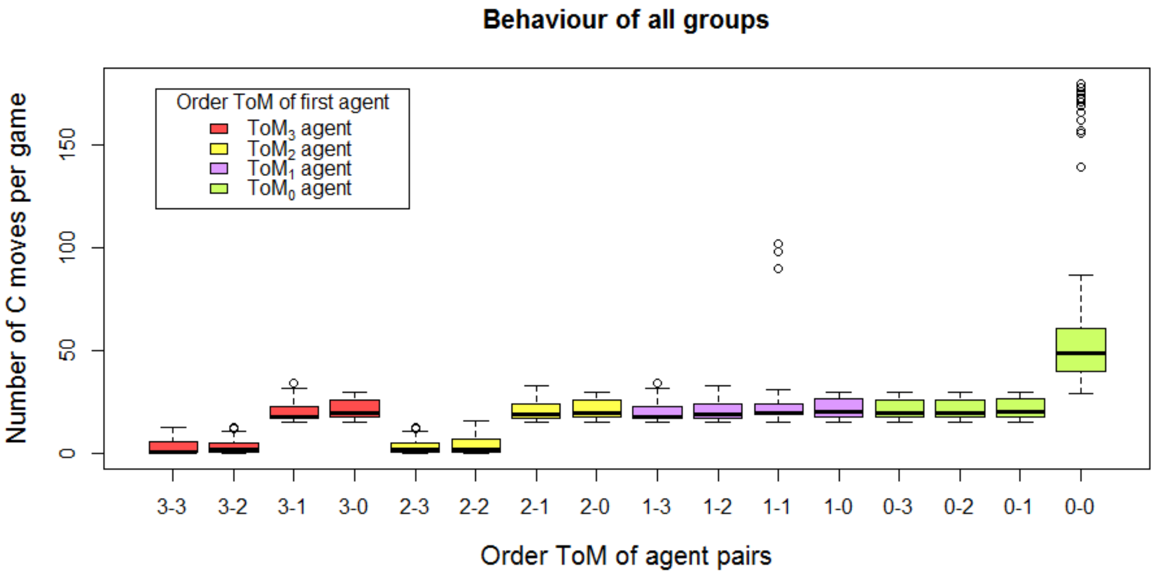

Figure 3.1: Every box shows the number of cooperative moves that were made during the 1000 games that the agent pair played. The position of the agents are interchangeable, so the data for an x-y pair is the same as a y-x pair.

ence the same amount of cooperation and betrayal. More cooperation in the memory could be evidence that there may be a threshold, where if there is enough cooperation in their past the agents will fall into a stable C-C situation. If the memory where to be mirrored it would explain why the observed behaviour of both agents was the same. However, the memory was not mirrored or even particularly similar from both agents' perspectives and there was no more cooperation than in other games. It is therefore unknown why the agents' moves became synchronized with such high cooperation after the forced defection period.

Since the data were non-normally distributed the Kruskal-Wallis test was used. This test indicated that there was a different amount of cooperation in at least one of the groups (df = 9, P < 0.001).

A Nemenyi post-hoc test revealed that the 0-0 agent pair had significantly more cooperative moves than any of the other groups. The pairs 2-0, 3-0, 1-0, 1-1 did not differ significantly from each other. The pairs 2-1 and 3-1 also did not significantly differ from each other. The pairs 3-3, 3-2 and 2-2 also did not show any significant difference from each other. Which groups differ significantly from each other has been portrayed in Table 3.1. In which the x and y-axis both show the order ToM an agent was capable of using. For example, the top-left cell represents the 0-0 agent pair, while the bottom-right cell represents the 3-3 agent pair.

The number of cooperative moves is highest in the pink cell, which only contains the 0-0 agent pair. Which is followed by the red cells, then the orange cells and the lowest number of cooperative moves is found in the yellow cells.

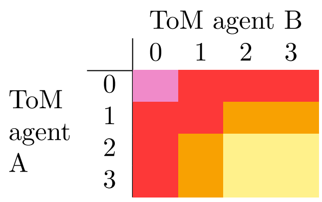

Table 3.1: When the color between two cells is different it indicates that there is a significant difference in the amount of cooperation between these agent pairs.

4 Discussion and Conclusions

From our results we observe that the agent-pairs with the lowest orders of ToM had the highest number of cooperative moves. In the introduction of this paper we asked whether it is theory of mind that leads to cooperation and, more specifically, whether a higher order theory of mind will get an agent out of a D-D situation and into a C-C situation. We expected that we would see more cooperation in agents with higher levels theory of mind. When we compare our hypothesis with our results we see that the trend we had expected cannot be found in the agents' behaviour. In fact, the trend that we observe in our results is the opposite of what we had expected. To answer our aforementioned questions, theory of mind on its own does not lead to more cooperation. In our results higher orders of theory of mind actually lead to more defection. There are several instances to be found in the literature that support these findings.

In earlier research it has been observed that in single-shot games human subjects with higher orders of ToM were less likely to cooperate (DeAngelo and McCannon, 2017).

Theory of mind not leading to more cooperation can also be observed in the natural world. Although theory of mind is often associated with social skills, there are instances where people are capable of using theory of mind while lacking certain other social processes. People with psychopathic personality disorder, who are characterized by impaired empathy and egotistical traits (Skeem, Polaschek, Patrick, and Lilienfeld, 2011), are capable of using normal levels of ToM. Psychopaths have been observed to defect more than people from the general population when playing the prisoner's dilemma and in real life social situations (Mokros, Menner, Eisenbarth, Alpers, Lange, and Osterheider, 2008).

The opposite has also been observed, that theory of mind is not always needed in cooperative behaviour. This can be seen in people who lack higher levels theory of mind, such as young children, and in the behaviour of animals that we assume do not have theory of mind. These people and animals are capable of cooperative behaviour even though they have very little to no theory of mind (Tomasello, 2009).

In conclusion, we found that on its own theory of mind does not lead to more cooperation in the prisoner's dilemma and, in fact, leads to more defection.

4.1 Future research

4.1.1 Discount Factor

The discount factor influences how important the gains from future rounds are to the agent; it shows how willing an agent is to lose this round to win a future one. Perhaps it would be interesting to link the discount factor value to how much confidence an agent has in its ability to predict its opponent's moves. When you have low confidence that you can predict your opponent accurately, it makes sense to not want to rely on future gains, but when you have very high confidence that you have accurately modeled your opponent those future wins could be just as important as you current one.

4.1.2 Opponent's strategy

Our agents rely on simulating their opponent and assume that their opponent acts and thinks like them, but what if that is not the case? A famous strategy in the prisoner's dilemma is tit-for-tat, in which an agent responds to its opponent by playing the same move its opponent played last round. If an opponent were to use that strategy against one of our agents, the agent would be unable to accurately predict its opponent, no matter what order ToM strategy it uses, since the opponent doesn't function like our agent. Instead of assuming that the opponent works the exact same way as the agent, it would be an option to try and fit known strategies to the opponent. Then an agent would try to see whether their opponent is playing by the tit-for-tat strategy, for example, and simulate an agent that uses that strategy to predict its opponent's moves.

4.1.3 Prospective thinking

The agents in this experiment were capable of prospective thinking, by thinking ahead about how their actions would influence the next three rounds of the game. We kept the amount of thinking steps static, but it would be interesting to see how the amount of thinking steps would influence the behaviour of the agents.

References

Luca Barlassina and Robert M. Gordon. Folk psychology as mental simulation. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2017 edition, 2017.

Colin F Camerer, Teck-Hua Ho, and Juin Kuan Chong. Behavioural game theory: Thinking, learning and teaching. In *Advances in Understanding Strategic Behaviour*, pages 120–180. Springer, 2004.

Gregory DeAngelo and Bryan C. McCannon. Theory of mind predicts cooperative behavior. *Economics Letters*, 155:1–4, 2017.

Adam S Goodie, Prashant Doshi, and Diana L Young. Levels of theory-of-mind reasoning in competitive games. *Journal of Behavioral Decision Making*, 25(1):95–108, 2012.

Peter Kinderman, Robin Dunbar, and Richard P Bentall. Theory-of-mind deficits and causal attributions. *British Journal of Psychology*, 89(2): 191–204, 1998.

Steven Kuhn. Prisoner's dilemma. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2017 edition, 2017.

Andreas Mokros, Birgit Menner, Hedwig Eisenbarth, Georg W. Alpers, Klaus W. Lange, and Michael Osterheider. Diminished cooperativeness of psychopaths in a prisoner's dilemma game yields higher rewards. *Journal of Abnormal Psychology*, 117(2):406–413, 2008. doi: 10.1037/0021-843X.117.2.406. ID: 2008-05639-015.

John Monterosso, George Ainslie, P. A. C Pamela Toppi Mullen, and Barbara Gault. The fragility of cooperation: A false feedback study of a sequential iterated prisoner's dilemma. *Journal of Economic Psychology*, 23(4):437–448, August 2002. ID: 271667.

David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(04):515–526, 1978.

Jennifer L. Skeem, Devon L. L. Polaschek, Christopher J. Patrick, and Scott O. Lilienfeld. Psychopathic personality. *Psychological Science in the Public Interest*, 12(3):95–162, 2011. doi: 10.1177/1529100611426706. PMID: 26167886.

Michael Tomasello. *Why we cooperate*. MIT press, 2009.

Michael Tomasello, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. In search of the uniquely human. *Behavioral and brain sciences*, 28(5):721–727, 2005.