

2015

Predicting Trust Dynamics and Transfer of Learning in Games of Strategic Interaction as a Function of a Player's Strategy and Level of Trustworthiness

Michael Collins

Ion Juvina

Wright State University - Main Campus, ion.juvina@wright.edu

Gary R. Douglas

Wright State University - Main Campus, douglas.45@wright.edu

Kevin A. Gluck

Follow this and additional works at: <https://corescholar.libraries.wright.edu/psychology>



Part of the [Psychology Commons](#)

Repository Citation

Collins, M., Juvina, I., Douglas, G. R., & Gluck, K. A. (2015). Predicting Trust Dynamics and Transfer of Learning in Games of Strategic Interaction as a Function of a Player's Strategy and Level of Trustworthiness. *24th Conference on Behavior Representation in Modeling and Simulation*.

<https://corescholar.libraries.wright.edu/psychology/521>

This Conference Proceeding is brought to you for free and open access by the Psychology at CORE Scholar. It has been accepted for inclusion in Psychology Faculty Publications by an authorized administrator of CORE Scholar. For more information, please contact corescholar@www.libraries.wright.edu, library-corescholar@wright.edu.

Predicting Trust Dynamics and Transfer of Learning in Games of Strategic Interaction as a Function of a Player's Strategy and Level of Trustworthiness

Michael Collins^{a, b}

Michael.Collins.74.ctr@us.af.mil
937-694-1065

Ion Juvina^a

Ion.juvina@wright.edu

Gary Douglas^a

Douglas.45@wright.edu

Kevin Gluck^b

Kevin.Gluck@us.af.mil

^a Department of Psychology, Wright State University,
3640 Colonel Glenn Hwy, Dayton, OH 45435

^b Air Force Research Laboratory, 2620 Q Street, Building, 852
Wright-Patterson Air Force Base, OH 45433

Keywords:

Cognitive Modeling, A Priori Model Prediction, Strategic Interaction, Trust Dynamics, and Transfer of Learning

ABSTRACT: *Individuals playing a sequence of different games have shown to learn about the other player's behavior during their initial interaction and apply this knowledge when playing another game with the same individual in the future. Here we use a published computational cognitive model to generate predictions for an upcoming human study. The model plays both Prisoner's Dilemma and Chicken Game with a confederate agent who uses one of two predetermined strategies and whose level of trustworthiness is manipulated. We go beyond the standard postdictive practice and adopt the increasingly popular practice of using the model to make a priori predictions before the human data will be collected in an upcoming study.*

1. Introduction and Background

How people learn to trust one another over time and how they use this information to inform their future decisions is a question relevant to many aspects of human interaction. Trust is defined as “the willingness of a party to be vulnerable to the actions of another party based on the expectations that they will perform a particular action” (Mayer, Davis, & Schoorman, 1995). Trust relationships have been proposed to be self-sustaining once developed, allowing individuals to forgo re-evaluation of a person after they have been determined to be trustworthy (Hardin, 2002). Yamagishi, Kanazawa, Mashima, and Terai (2005) found that when participants played a modified version of the game Prisoner's Dilemma (PD), where participants could choose the amount of points they could risk during each round, over time participants would gradually increase the number of points they would risk as the individuals began to establish trust for one another. Consistent with these results, Castelfranchi & Falcone (2010) suggest that trust mitigates risk and develops through gradual risk-taking between two individuals.

In order to study how individuals behave in different situations, both economic and psychological research have

used games of strategic interaction. A game represents an abstraction of a real-world scenario in which participants can win and lose points based on the behavior of both players. Participants can play either with another human participant (e.g., Juvina, Saleem, Martin, Gonzalez, & Lebiere, 2013) or with a preprogrammed strategy (e.g., Juvina, Lebiere, Martin, & Gonzalez, 2012).

Two different strategies that have been used in place of human participants during these games of strategic interaction are the Tit-for-Tat (T4T) (Axelrod, 1984) and the Pavlov-Tit-for-Tat strategy (PT4T) introduced by Juvina, Lebiere, Gonzalez and Saleem (2012). T4T is a simple strategy, which repeats on round N the same choice that the other player made on round N-1. The PT4T strategy is a combination of two different strategies, T4T and Pavlov. Pavlov is another simple strategy that continues to choose the same choice on round N as long as it earned points with that choice on round N-1, only changing choices on round N when it lost points on round N-1. The PT4T strategy repeats the other player's move from N-1 on round N, just as the T4T strategy, except for when the strategy and the other player make opposite choices and the strategy earns points on that round. Instead of switching to the other player's choice as the

T4T strategy would, the PT4T strategy repeats its previous choice, as the Pavlov strategy would. The PT4T strategy was created based on analysis of the repetition propensities (the probability to repeat a move following a certain outcome) of humans in PD and in an attempt to develop a strategy that had similar repetition propensities as humans (Juvina et al., 2012).

Previous research has found that when individuals play games of strategic interaction sequentially, they use the information gained about the other player from a previous game to inform their choices when playing with that person again (Juvina et al., 2013). Different explanations have been offered for why these transfer effects occur, such as a similarity between the games, the expectation of the other player to behave as they did in the past, or a strategy that was used during a simpler game continuing to be used in a more complex game (Knez & Cramer, 2000; Devetag, 2003; Bednar, 2012).

Juvina et al. (2013) found that these explanations failed to account for the transfer effects seen when repeated rounds of the games Prisoner’s Dilemma (PD) and Chicken Game (CG) were played sequentially. As an alternative explanation for why transfer effects occur between these games, Juvina et al. (2013) proposed that it is the increase in reciprocal trust between the two players that results in a transfer of learning occurring between these games, allowing them to find the optimal outcome faster in the second game compared to the first. Juvina, Lebiere, and Gonzalez (2014) implemented this idea of reciprocal trust in a computational cognitive model that replicates the transfer effects seen when the games PD and CG are played sequentially in either order.

The results in Juvina et al. (2014) were obtained by fitting the model post-hoc to the human data from Juvina et al. (2013) by manipulating certain model parameters. However, fitting the model post-hoc to a specific dataset does not ensure its validity and generalizability. In order to fit the human data, the model played against itself, using both the same parameters and learning mechanisms to determine how to play both games. This is problematic when trying to understand real world scenarios where individuals are likely to have different goals and understandings of the current situation. Due to these differences, it has not yet been shown that the model can account for human behavior when playing against an individual who has a different approach and a different level of trustworthiness.

We are attempting to validate the model used in Juvina et al. (2014) by using the model to simulate the results of an upcoming study to be conducted with human participants. The model will play two games sequentially, either PD and CG in varying orders or one of the two games twice with a

preprogrammed confederate agent. The confederate agent will use one of two predetermined strategies and will have varying levels of trustworthiness. A comparison of the model’s predictions to the behavior of human participants will allow for an opportunity to examine in what types of situations the model can predict the behavior of human participants. In this article, a brief overview of the model and the experimental design of the simulation is offered, along with a discussion of the model’s predictions for the upcoming study to be conducted with human participants.

1.1 The Games

Participants will play repeated rounds of the same two games used in Juvina et al.’s (2013) original study, which are PD and CG. Both PD and CG are mixed motive non-zero sum games and are represented by their own payoff matrix (Fig 1.1). During each round in a game, both Player 1 (P1) and Player 2 (P2) choose to either defect (A) or cooperate (B). Based on the choices made by both players during every round, P1 or P2 either win or lose a certain number of points.

			Prisoner's Dilemma Payoff Matrix					Chicken Payoff Matrix	
			A	B				A	B
P1	A	-1	-1	4	-4	-4	-4	4	-1
	B	-4	4	1	1	-1	4	1	1

Fig 1.1. The payoff matrix for the game Prisoners dilemma (left) and Chicken Game (right).

When either PD or CG is played continually and both players do not know how long they will play, each game has a different optimal outcome. In PD, the optimal outcome over the course of the game is for both players to choose B (mutual cooperation) in order to earn one point each during each round (Fig 1.1). In CG, the optimal outcome is for both players to asymmetrically alternate between choosing A and B, earning three points every other round (Fig 1.1). However, when playing either CG or PD, attempting to choose the optimal outcome is risky. If only one player understands the benefits of sustaining the mutual cooperation or alternation outcome and is willing to reciprocate, then the player who attempts the optimal strategy will lose points as the other player gains points. To avoid this, players must learn to mutually cooperate with one another by sustaining the optimal outcome throughout the game, which maximizes their payoffs when either PD or CG is played repeatedly (Juvina et al., 2013). Due to the fact that each game has a different optimal outcome, the behavior of both players should change along with the games that are played.

Although PD and CG have different payoff matrices, certain characteristics are similar across both games. There are both surface and deep similarities. The surface similarity between PD and CG that is relevant in this context is that both players during either game can choose B to earn one point during each round. Both games also share a deep similarity that is both players mutually cooperating with each another brings about the optimal outcome when either game is played repeatedly. Players can mutually cooperate by both choosing B in PD and asymmetrically alternating between A and B in CG (Juvina et al., 2013). Juvina et al. (2013) has found that when PD and CG are played sequentially the transfer effects between these games occur along both the surface and deep similarities. In particular, more mutual cooperation was seen in PD when played after CG and more alternation was seen in CG when played after PD.

1.2 The Model

A brief overview of the model used to generate the predictions of the upcoming study is given here; a more detailed description of the model can be found in Juvina et al. (2014). The model was built in ACT-R (Adaptive Control of Thought - Rational), which is both a cognitive architecture and a theory of human cognition (Anderson, 2007). In ACT-R, different modules interact with each other in order to complete a task. In the model used for this study, two memory modules are used in order to play both games; these are the declarative and procedural modules. The declarative module stores information that the model has learned from the environment. The procedural memory allows for action selection reinforced through reward patterns that occur within the environment (Anderson, 2007). Both modules are used together to account for human behavior in the two games when played independently and sequentially.

In order for the model to be able to play either game, it needs to be aware of the interdependence between itself and the other player; to do this the model uses instance-based learning (IBL: Gonzalez, Lerch, & Lebiere, 2003). In IBL, past instances of an event are stored in a model's declarative memory to be recalled later, and inform future decisions. When the model is in a situation similar to a previous experience, it uses information stored in its declarative memory to make a decision about what to do in its current situation. At each round, the model stores in its declarative memory the previous move of both itself and the other player along with the other player's move for the current round. Throughout both games, each time the model stores a copy of a previous instance that has already been placed in its declarative memory it increases the probability that that specific instance will be recalled when placed in a similar situation again, as controlled by ACT-R's activation equations (Anderson, 2007).

To account for the behavior of the human participants in each game, the model uses both IBL and reinforcement learning. During each round, the model attempts to recall a previous instance from memory using both its own and the other player's previous move as retrieval cues. The stored previous instances in the model's declarative memory allow it to recall what the other player's next move was when placed in that situation before. The model predicts that the other player will choose the move that was chosen more frequently when placed in similar situations in the past. The model then chooses to cooperate or defect depending on which choice has the greatest utility given the model's prediction of the other player's move. Previous rewards the model has received for cooperating and defecting in similar contexts (i.e., the other players expected next move based on the previous move of the other player and the model) determine the utility or the value of these choices to the model (Juvina et al., 2014).

In order to account for the deep transfer effects seen when PD and CG are played sequentially, two trust accumulators and three different reward functions were added to the model. The two accumulators are called trust and trust-invest. Each accumulator starts at zero at the beginning of the first game and increases or decreases depending on the moves both the model and the other player make after each round. The trust accumulator increases when both players either mutually cooperate or when the model defects and the other player cooperates. It decreases when both players mutually defect or when the model cooperates and when the other player defects. The trust-invest accumulator increases with mutual defections and decreases with unreciprocated cooperation. Throughout either game the current levels of the trust and trust-invest accumulators determine the model's current reward function.

Three reward functions are used which reinforce the model's choices differently for each of the four possible outcomes that can occur during a game, in turn affecting the model's behavior. By alternating between three different reward functions, the model uses the reward function that is most applicable to its current situation. The reward function that is applied to the current round of the game is determined by the level of the trust and trust-invest accumulator. When the trust accumulator is positive, the model is reinforced for increasing the payoff of both players. When only the trust-invest accumulator is positive, the model is reinforced for increasing the payoff of the other player. When both accumulators are at or below zero, the model is reinforced for maximizing its own payoff and minimizing the payoff of the other player.

2. The Experiment

The model predictions presented in this paper were generated by simulating a fully balanced 4 x 2 x 2

experiment that will be conducted with human participants. Participants will play both PD and CG or one of these two games twice. Instead of participants playing games with one another as in Juvina et al.'s (2013), participants will play with a "confederate agent", implemented as a software agent. The confederate agent will use one of two predetermined strategies and the trustworthiness of the agent will be controlled, while playing both games. The model was run in conditions identical to those that future participants will be placed in.

On Qualtrics.com, the online platform that will be used to run the upcoming experiment, we created sixteen conditions with each possible combination of game order, confederate agent's strategy, and trustworthiness. In each condition, ten preprogrammed versions of each game were developed to ensure random variability in the behavior of the confederate agent. Once the experiment begins participants will first be randomly assigned to a condition and then randomly assigned to play one of the ten possible versions of each of the two games they will play during the experiment. The experimental protocol for the upcoming study was copied when generating model predictions, simulating fifty human participants in each condition.

2.1 The Confederate Agent

The confederate agent will utilize one of two predetermined strategies throughout both games. The T4T strategy will choose on round N the choice that the other player made on round N-1. The PT4T strategy will reciprocate mutual cooperation and defection, but will not reciprocate unilateral cooperation.

Along with using one of two predetermined strategies, the confederate agent's trustworthiness will be manipulated and randomness will be added into its behavior. To accomplish this, the confederate agent will either cooperate or defect a certain number of times throughout each game at random times. In the high trustworthiness (HT) condition the confederate agent will cooperate and in the low trustworthiness (LT) conditions the confederate agent will defect. For this experiment, we wanted to create a confederate agent that would generate significant differences in the outcomes that were chosen across all conditions. To accomplish this, multiple model predictions for all conditions were run by varying the number of rounds the confederate agent employed its strategy (reactive strategy – T4T or PT4T) and automatic cooperation or defection (fixed strategy). We found that, because PT4T is inherently less trustworthy than T4T (i.e., more apt to defect), to avoid the model only predicting a high frequency of mutual defection during the PT4T HT conditions, a larger percentage of cooperation was needed to raise the strategies trustworthiness. For this experiment, during the T4T conditions, the confederate agent will

employ the T4T strategy randomly during 90% of the game, while randomly employing its fixed strategy during 10% of the game. During the PT4T conditions, the confederate agent will employ the PT4T strategy randomly during 65% of the game and randomly employ its fixed strategy during 35% of the game.

3. Results and Discussion of the Model's Predictions

We computed the frequency of five relevant outcomes during each round in every condition over the fifty different model runs. In order to determine instances of asymmetrical alternation, rounds where one player chose to defect while the other player cooperated or vice versa on round N and had both chosen the opposite choices on round N-1 were identified. The frequency of alternation during each round across all conditions was computed like all other outcomes. Because of the limitation of space in this paper, we cannot report all of the results. All of the model's predictions are available for viewing and can be downloaded at (<http://psych-scholar.wright.edu/ijuvina/publications>). A linear mixed effects analysis (LME) was used to assess the effect of strategy, trustworthiness of the confederate agent, and order in which the games were played on the predicted frequency of the five outcomes. P-values were obtained using a likelihood ratio test comparing a full to a reduced model. The 95% confidence intervals for the effects predicted by the LME are also reported. It should be noted that the confidence intervals that are reported are large, which is expected given the large variability generated by each ACT-R model, the randomness added to the confederate agent, and the multitude of experimental conditions. The model's predictions will be compared to human data from each condition, once the experiment has been run.

Transfer effects were assessed using a paired t-test, run on the frequency of each outcome during the first game compared to the frequency of that outcome when the same game was played second against a confederate agent of the same strategy and level of trustworthiness. Significant results indicate that the order in which the model played the game affected the frequency that an outcome was chosen during that game.

3.1 Effects of Trustworthiness

One clear difference seen across the high and low trustworthiness conditions in the model's predictions is the level of the trust accumulator. A t-test run on the round-by-round average of the magnitude of the trust accumulator across the simulated low ($M = -66.86$, $SD = 38.11$) and high ($M = 62.36$, $SD = 39.17$) trustworthiness conditions was

found to be significant ($t(49) = 66.87, p < .001$). The model's current level of the trust accumulator affects which current reward function is used and will determine whether the model will attempt to maximize its own payoff or the payoff of both players. The difference in the trust accumulator between the simulation of the high and low trustworthiness conditions indicates that the experimental manipulations of trustworthiness were effective. Based on its level of trust, the model predicts a difference in the frequency that mutual defection will occur in both games, despite differences in the strategy used by the confederate agent and order.

A LME was run with the average predicted frequency of mutual defection as a dependent variable, trustworthiness of the confederate agent as a fixed effect, with strategy, order, and round as random factors. A likelihood ratio test was run and found that the trustworthiness of the confederate agent was found to have a significant effect on the predicted frequency of mutual defection ($X^2(1) = 277.3, p < .001$), increasing the frequency of mutual defection by $75.07\% \pm .6\%$, 95% CI [52% , 98 %], during the simulated low trustworthiness conditions compared to $15.4\% \pm 6.5\%$, 95% CI [0 , 37.33%], in the simulated high trustworthiness conditions (Fig. 1.2).




Fig 1.2. The average round-by-round frequency that mutual defection was chosen across all of the simulated high (dashed red line) and low (solid black line) trustworthiness conditions.

The trustworthiness of the confederate agent determines whether it will cooperate (high trustworthiness) or defect (low trustworthiness) for a specific number of times (10% of the rounds in the T4T and 35% of the rounds in the PT4T) over the course of the game at random times. The model predicts that participants will be sensitive to the trustworthiness of the confederate agent, responding by defecting more throughout the low trustworthiness conditions and less during the high trustworthiness conditions.

3.2 Effects of Strategy

The two types of strategies used by the confederate agent have different criteria for deciding what choice to choose during each round; these differences limit how quickly the model can change from one outcome to another and the

outcomes that can be achieved during a game. For example, continual alternation is an outcome that can only be achieved with the T4T strategy and not with the PT4T strategy. Continual mutual cooperation is also an outcome that is harder to achieve with the PT4T strategy, because it is inherently less trustworthy (i.e., more apt to defect). It is the differences in the behavior of these two strategies used by the confederate agent that affected the predicted frequency in which the optimal outcomes will be chosen despite differences in the trustworthiness of the confederate agent or the order in which the games are played.

A LME was run with the average predicted frequency of mutual cooperation as a dependent variable, strategy as a fixed factor, with trustworthiness, order, and round as mixed effects. A likelihood ratio test was conducted and found that the strategy implemented by the confederate agent significantly affected the predicted frequency of mutual cooperation ($X^2(1) = 68.867, p < .001$). The T4T strategy had a larger affect on the predicted frequency of mutual cooperation, increasing its predicted frequency by $25.1\% \pm .7\%$, 95% CI [0 , 70%] compared to when the confederate agent used the PT4T strategy, increasing the predicted frequency of mutual cooperation by only $19\% \pm 13.7\%$, 95% CI [0 , 62%] (Fig 1.3). A second LME was run with the average predicted frequency of alternation as a dependent variable, strategy as a fixed factor, with trustworthiness, order, and round as random factors. Similar to mutual cooperation, the strategy used by the confederate agent was found to have a significant effect on the predicted frequency of alternation ($X^2(1) = 392.21, p < .001$). Conditions where the confederate agent used the T4T strategy had a larger affect on the predicted frequency of alternation, increasing the frequency by $12.9\% \pm 0.4\%$, 95% CI [6% , 30%] in conditions where the confederate agent used the T4T strategy compared to only $4\% \pm 6\%$, 95% CI [0% , 20%] when it used the PT4T strategy (Fig 1.3).

The strategy used by the confederate agent was also found to have a significant effect on the predicted frequency of mutual defection, controlling for trustworthiness and order. A LME was run with the average predicted frequency of mutual defection as a dependent variable and strategy as a fixed factor, with trustworthiness, order, and round as random effects. A likelihood ratio test was conducted and found that the strategy used by the confederate agent had a significant effect on the predicted frequency of mutual defection ($X^2(1) = 574.02, p < .001$). Conditions where the confederate agent used the PT4T




Fig 1.3. A comparison of the average predicted frequency per round of three different outcomes: mutual cooperation (CC), alternation (ALT), and mutual defection (DD), across all the Tit-for-Tat (T4T, solid black line) and Pavlov-Tit-For-Tat (PT4T, dashed red line) conditions. The 95% confidence intervals per round for each outcome and condition are also plotted

strategy had a larger effect on the predicted frequency of mutual defection, increasing its frequency by $54.1\% \pm 29\%$, 95% CI [0%, 100%], compared to when the confederate agent used the T4T strategy increasing the predicted frequency of mutual defection by only $36.31\% \pm .6\%$, 95% CI [0% , 100%] (Fig 1.3).

The model predicts that participants will react differently to the two different strategies used by the confederate agent. Alternation and mutual cooperation are both predicted to occur at a higher frequency during all of the T4T conditions compared to the PT4T conditions. A higher predicted frequency of alternation occurring during the T4T conditions would be expected, because the PT4T strategy cannot continually alternate throughout the game like the T4T strategy. However, the T4T and PT4T strategy can both mutually cooperate throughout a game. The difference that the frequency of mutual cooperation is predicted to occur is caused by the strategies' behavior during the experiment when played with repeatedly, because repeated instances of mutual cooperation are harder to obtain with the PT4T strategy than with the T4T strategy. In addition, as is seen in the model's predictions, the PT4T condition is predicted to have a higher frequency of mutual defection across all conditions, which would affect the model's trust in the confederate agent, leading it to cooperate less in conditions where the confederate agent used the PT4T strategy compared to the T4T strategy.

3.3 Effects of Order

The optimal outcomes that are chosen during the experiment depend on the games that are played during each condition. For example, alternation is the optimal

outcome in CG, but is not an optimal outcome in PD, because alternating between a payoff of +4 and -4 points per round leads to a net gain of 0 for both players. While playing PD, mutual cooperation is the optimal strategy and though mutual cooperation is a possible outcome in CG, it leads to a sub-optimal outcome compared to alternation, +1 point per round compared to +3 points every other round. Juvina et al. (2013) found that order also affects the frequency of the optimal outcomes during a game. The optimal outcome in either PD or CG occurred more frequently when it was played after the other game compared to when played first. Due to the effects that order has been seen to have on the outcomes that are chosen over the course of both games, the model will predict a significant difference in the frequency of the two optimal outcomes over the course of the two games depending on the order that they are played.

A LME was run with the average predicted frequency of mutual cooperation as a dependent variable, order as a fixed effect, with trustworthiness and strategy of the confederate agent and round as random effects. A likelihood ratio test was conducted and found that the order in which the games were played in a condition significantly affected the frequency of mutual cooperation ($X^2(3) = 712.98$, $p < .001$), increasing the predicted frequency of mutual cooperation by $36.6\% \pm 1\%$, 95% CI [0% , 79%] in the simulated conditions when PD was played repeatedly (PDPD order), $28.47\% \pm 1\%$, 95% CI [0% , 71%], when PD was played before CG (PDCG order), $13.10\% \pm 1\%$, 95% CI [0% , 71%], when CG was played before PD (CGPD order), and $10\% \pm 12.3\%$, 95% CI [0% , 51%], when CG was played twice (CGCG order).

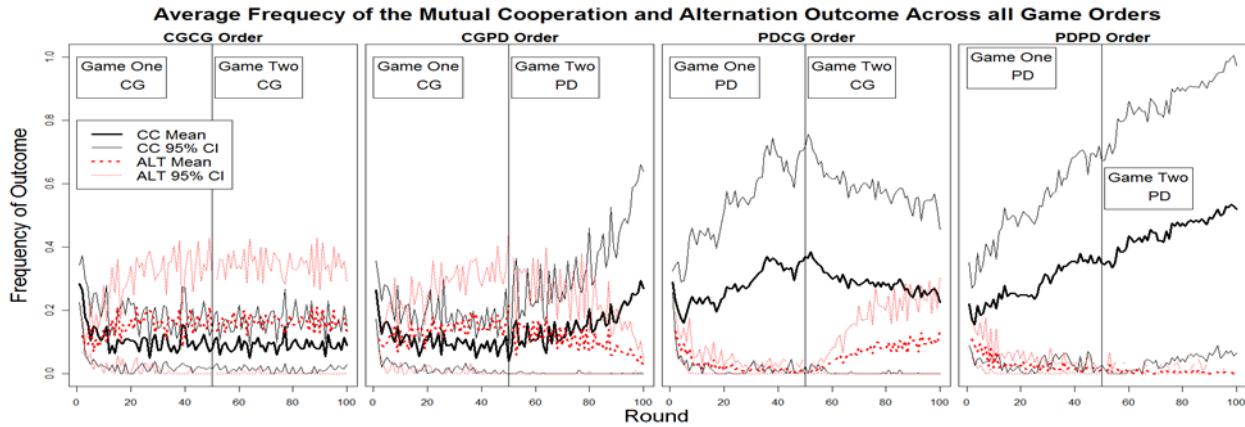


Fig 1.4. Comparison of the model predictions of the average predicted frequency per round of two different optimal outcomes: mutual cooperation (CC, solid black line) and alternation (ALT, dashed red line), across all of the different orders that PD and CG were played in. The 95% confidence intervals per round for each outcome and condition are also plotted.

To test the significance of the effect of order on the predicted frequency of alternation, a LME was run with the average predicted frequency of alternation as a dependent variable, order as a fixed factor, with trustworthiness and strategy of the confederate agent and round as random effects. The order in which the games were played was found to significantly affect the predicted frequency of alternation, opposite that of the predicted frequency of mutual cooperation ($X^2(3) = 712.98, p < .001$). Game order affected the frequency of alternation by $15.5\% \pm 5.9\%$, 95% CI [0% , 33%], in simulated conditions with the CGCG order, $11.9\% \pm .5\%$, 95% CI [0% , 23%], in the CGPD order, $4.95\% \pm .05\%$, 95% CI [0% , 23%], in the PDCG order, and $1.86\% \pm .05\%$, 95% CI [0% , 20%], in the PDPD order (Fig 1.4).

The affect that the order games were played had on the predicted frequency of the optimal outcomes show that in conditions where the same game is played repeatedly, such as in the PDPD and CGCG order, the model predicts that the frequency of the optimal outcome for that game will continue to increase throughout the condition. The model also makes an uncharacteristic prediction about the frequency that mutual cooperation and alternation in the conditions simulated with the PDCG and CGPD order. It would be expected based on results from Juvina et al. (2013), that conditions with the PDCG order would have a higher frequency of alternation than the CGPD order, and that the CGPD order would have a higher frequency of mutual cooperation than with the PDCG order. Instead, the model predicts that when PD and CG are played in sequence, the highest frequency of mutual cooperation will be in conditions with the PDCG order and the highest frequency of alternation will occur in conditions with the CGPD order.

3.4 Predicted Transfer Effects

Previous results with human pairs have found that when PD and CG were played in sequence, transfer effects between these two games occur along both their surface and deep similarities (Juvina et al., 2013). The same transfer effects have also been found when cognitive models were paired with one another (Juvina et al., 2014). In contrast, when a cognitive model was paired with a pre-programmed agent as in the current study, no deep transfer effects are predicted; the model only predicts surface transfer effects. Mutual cooperation in the T4T HT condition is predicted to occur at a higher frequency during CG when played after PD compared to when played before PD ($t(49) = -21.8871, p < .001$). The same prediction about the frequency of mutual cooperation is made during the PT4T HT condition. Mutual cooperation is predicted to occur at a higher frequency during CG when played after PD compared to when played before PD ($t(49) = -38.429, p < .001$).

The surface transfer effect of mutual cooperation in the PDCG order during the PT4T HT condition is amplified by the limitations of the confederate agent's strategy. Because continual alternation cannot be achieved with the PT4T strategy, mutual cooperation, a sub-optimal outcome in CG, is left as the only satisfactory outcome that can be achieved given the behavior of the confederate agent. One possible explanation for the lack of deep transfer effects in the model's predictions is the difference between the behavior of the confederate agent and an actual human player. The confederate agent is simpler than the model (even with the added randomness) and does not learn from the interaction with the model throughout the game. If confirmed, the prediction of a lack of deep transfer will strengthen the claim made in Juvina et al. (2013, 2014) that joint learning and reciprocal trust are key ingredients for a deep transfer of learning in games of strategic interaction.

4. Conclusion

In summary, we are validating a computational cognitive model that has shown to be able to account for the transfer effects that are observed when the games PD and CG are played repeatedly and in sequence with human participants. In order to validate the model, we have made a priori model predictions about the behavior of human participants when playing against a preprogrammed confederate agent across a variety of conditions. From the model's predictions we have developed five hypotheses for the upcoming study.

H1: We predict that mutual defection will be chosen more across all of the low trustworthiness conditions compared to the high trustworthiness conditions.

H2: We predict both optimal outcomes (i.e., mutual cooperation and alternation) will be chosen at a higher frequency in conditions where the confederate agent uses the T4T compared to the PT4T strategy.

H3: We predict that the frequency of both optimal outcomes (i.e., mutual cooperation and alternation) will depend on the order that games are played in a condition.

H4: We predict that across the sixteen conditions no deep transfer of learning will occur.

H5: We predict that across the sixteen conditions surface transfers of learning will only occur with the mutual cooperation outcome in the PDCG PT4T HT and PDCG T4T HT condition.

We expect to run the study in 2015. A subsequent publication will reveal the actual empirical results and degree of model predictive validity.

5. References

- Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* New York: Oxford University Press.
- Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Bednar, J., Chen, Y., Xiao Liu, T., & Page, S. E. (2012). Behavioral spillovers and cognitive load in multiple games: An experimental study. *Games and Economic Behavior*, 74(1), 12–31.
- Castelfranchi, C., & Falcone, R. (2010). *Trust Theory: A Socio-Cognitive and Computational Model*. Chichester: John Wiley and Sons.
- Devetag, G. (2005). Precedent transfer in coordination games: An experiment. *Economic Letters*, 89(2), 227–232.
- Gonzalez, C., Lerch, F. J., & Lebiere, C. (2003). Instance-based learning in real-time dynamic decision making. *Cognitive Science* 27 (4), 591-635.
- Hardin, R. (2002). *Trust and Trustworthiness*. New York: Russell Sage Foundation.

- Juvina, I., Lebiere, C., Martin, J., & Gonzalez, C. (2011). Intergroup Prisoner's Dilemma with Intragroup Power Dynamics. *Games*. 2(1), 21-51.
- Juvina, I., Lebiere, C., & Gonzalez, C. (In press, Nov. 2014). Modeling trust dynamics in strategic interaction. *Journal of Applied Research in Memory and Cognition*.
- Juvina, I., Lebiere, C., Gonzalez, C., & Saleem, M. (2012). Intergroup Prisoner's Dilemma with Intragroup Power Dynamics and Individual Power Drive. In *Social Computing, Behavioral-Cultural Modeling and Prediction*, 290-297.
- Juvina, I., Saleem, M., Martin, J. M., Gonzalez, C., & Lebiere, C. (2013). Reciprocal trust mediates deep transfer of learning between games of strategic interaction. *Organizational Behavior and Human Decision Processes*, 120(2), 206-215.
- Knez, M., & Camerer, C. (2000). Increasing cooperation in Prisoner's Dilemmas by establishing a precedent of efficiency in coordination games. *Organizational Behavior and Human Decision Processes*, 82(2), 194–216
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of management review*, 20(3), 709-734.
- Yamagishi, T., Kanazawa, S., Mashima, R., & Terai, S. (2005). Separating Trust from Cooperation in a Dynamic Relationship Prisoner's Dilemma with Variable. *Dependence, Rationality and Society*, 17(3), 275-308.

Author Biographies

Michael Collins is a current undergraduate in the Department of Psychology at Wright State University working as a research assistant at the Air Force Research laboratory. His research interest include computational cognitive modeling, game theory and social decision making.

Ion Juvina Ph.D. is an Assistant Professor in the Department of Psychology at Wright State University. His research interest span the areas of strategic decision making and executive control of cognition and affect. He teaches cognition and learning, psychological game theory, computational cognitive modeling, and experimental design.

Gary Douglas is a current human factors graduate student at Wright State University. His research interest include cognitive modeling, trust, and game theory.

Kevin Gluck Ph.D. is a Principal Cognitive Scientist at the Air Force Research laboratory and is currently the Acting Technical Advisor for the Warfighter Readiness Research Division.