Individual Differences and Levels of Analysis in Computational Models of Coordination

Alexander R. Hough (alexander.hough.1@us.af.mil)
Air Force Research Laboratory
Wright-Patterson AFB

Ion Juvina (ion.juvina@wright.edu)
Department of Psychology, College of Science and Mathematics, Wright State University
Dayton, OH

Abstract

Coordination failure is common and literature suggests individual differences play a role. Individuals are hypothesized to use strategies, form beliefs about others, and have different starting preferences. Extant data analysis and modeling efforts focus on average-level behaviors and often ignore individual differences in coordination strategies and their correspondence to cognitive mechanisms. Here, we leverage computational models to better understand individual differences and underlying cognitive mechanisms. We use experimental data from a coordination game to assess and compare a model from behavioral game theory, an extension of that model, and a recently developed cognitive model. This work presents challenges for modeling coordination dynamics, strategies, and how players form beliefs about others.

Keywords: Coordination; Group dynamics; Signaling; Coordination strategies; ACT-R; Cognitive model

Introduction

Coordination failure is common (Camerer, 2003; Riechmann & Weimann, 2008; Van Huyck, Battalio, & Beil, 1990, 1991), and could involve either miscoordination (i.e., failure to converge on a choice) or inefficiency (i.e. converging on a suboptimal choice). It is often attributed to the lack of salient focal points leading to efficient outcomes (Mehta, Starmer, & Sugden, 1994) and individual differences in starting strategies (Costa-Gomes, Crawford, & Iriberri, 2009; Van Huyck et al., 1990, 1991), persistence in trying to improve outcomes (Brandts, Cooper, & Weber, 2015), sensitivity to risks (Cachon & Camerer, 1996), and reciprocity (Offerman, 2002). To counteract these issues, individuals can form and update beliefs about what others will do (Camerer, 2003), engage in counterfactual thinking to consider what could have happened (Hough, O'Neill, & Juvina, 2021), and nudge others to make better choices by signaling or leading by example (Brandts et al., 2015; Hough et al., 2021). However, coordination dynamics and individual differences are not well understood. Extant research focuses on average behavior, rather than individual differences for analysis (Bortolotti, Devetag, & Ortmann, 2016; Leng, Friesen, Kalayci, & Man, 2018; Van Huyck et al., 1991, 1990) and modeling (Camerer & Ho, 1999; Costa-Gomes et al., 2009), which can lead to faulty conclusions about strategy use, strategy shifts due to learning, and differences between individuals (Siegler, 1987). To better understand coordination failure and individual differences, we turn to computational models that require detailed specification of mechanisms which serve as testable hypotheses and allow observation of "black box" processes. We use experimental data collected from the minimum effort game (i.e., MEG) (Van Huyck et al., 1990) to assess and compare a model from behavioral game theory, an extension of that model, and a recently developed cognitive model.

The MEG

The MEG is a weak link coordination game where each player chooses a level of effort between one and seven, and their payoff is determined by their choice and the group minimum (Table 1). There are seven coordination points

Table 1: MEG payoff matrix.

		Minimum Effort Choice in Group						
		1	2	3	4	5	6	7
Player Effort Choice	1	**70**						
	2	60	**80**					
	3	50	70	**90**				
	4	40	60	80	**100**			
	5	30	50	70	90	**110**		
	6	20	40	60	80	100	**120**	
	7	10	30	50	70	90	110	**130**

(i.e., Nash equilibria) represented diagonally in Table 1. Van Huyck et al. (1990, 1991) suggested players use risk and payoff dominant strategies. The risk dominant strategy is low risk, low reward, and is individually focused. Choosing one results in the same payoff regardless of other player's choices. The payoff dominant strategy is high risk, high reward, and is more group focused. The highest choice of seven can result in either the highest or the lowest payoff, depending on other's choices. These strategies serve as focal points, and over time, the minimum can become more influential. This was proposed as a simple explanation for the frequently observed negative trend in effort (i.e., efficiency), with the minimum and full (i.e., all player choices) feedback (Camerer, 2003; Leng et al., 2018; Van Huyck et al., 1991, 1990). However, there is evidence that players signal (Hough et al., 2021; Leng et al., 2018), engage in counterfactual thinking (Camerer & Ho, 1999; Hough et al., 2021), and speculation that individuals form beliefs about others "player type" corresponding to their initial preferences and future choices (Camerer, 2003).

MEG experiments (Bortolotti et al., 2016; Leng et al., 2018; Van Huyck et al., 1990, 1991) typically focus on average efficiency (i.e., minimum and average effort), and rarely explore group and individual behavior. However, Leng et al. (2018) identified signaling as alternating between the minimum and higher effort, Bortolotti et al. (2016) identified weak links as an early source of coordination failure, and Hough (2021); Hough et al. (2021) measured coordination by calculating intra-group variance and signaling by calculating each player's distance from the minimum. These contributions have not significantly improved our understanding of coordination dynamics or individual differences. In an attempt to address these issues, we turn to computational models. We discuss the experience-weighted attraction model (i.e., EWA) (Camerer & Ho, 1999) and an extended version to highlight limitations that motivated developing the cognitive model. We then compare model fits to data from a MEG experiment with four-human groups and full feedback (Hough, 2021)

Computational Models of the MEG

The EWA Model

EWA is based on forming and updating attractions towards all possible choices. It features initial choice attractions from prior experience, updating attractions weighted by recency and experience, and forgone payoffs that could have been earned (i.e., counterfactual thinking). The model has four parameters: 1) forgone payoff weight (δ) for counterfactual thinking, 2) past attraction decay (ϕ) and 3) experience decay (ρ) that control the growth rate of choice attractions, forgetting and recency effects, and 4) discrimination sensitivity (λ) to account for individual differences. The functioning of EWA can be described in three equations that calculate: experience, choice attractions, and choice probability. The experience equation calculates current experience (i.e., rounds), $N(t)$, based on previous experience, $N(t-1)$, that is depreciated, $\rho < 1$, each time there is a new experience (+1): $N(t) = \rho * N(t-1) + 1$. The choice attraction equation determines the attraction, A^j,for each choice, s^j, based on adding the depreciated previous attraction, $\phi * N(t-1) * A^j(t-1)$, to the current weighted payoff, $[\delta + (1-\delta) * I(s^j, s(t)))] * payoff$, and dividing it by current experience, $N(t)$. The current weighted payoff is equal to the payoff for the actual choice and is weighted by δ for forgone choices. Choice probability is determined by the Luce choice rule (Luce, 1956). Choice probability, $P^j(t+1)$, is based on a logistic transformation by raising the Euler's number, e, to the power of the choice attraction, $A^j(t)$, multiplied by the sensitivity parameter, λ, and is normalized by dividing it by the sum of all logistically transformed choice probabilities, $e^{\lambda * A^j(t)} / \Sigma_{k=1}^{m} e^{\lambda * A^k(t)}$.

Figure 1 shows EWA processes relating to the MEG. At the first round, EWA "knows" the payoff matrix, uses pregame experience and choice attractions to determine probabilities that serve as weights for sampling a choice. The model receives the minimum as feedback and calculates actual and weighted forgone payoffs. These payoffs are used to update experience and choice attractions, which are depreciated by experience, ρ, and attraction decay, ϕ. As the model only receives the minimum, it is not aware of all players choices or signaling efforts. To enable this capability, we extended EWA so it considers each player's choice as a hypothetical minimum. The extended EWA (i.e., EEWA) generates a set of attractions based on the minimum, then additional sets are generated for each hypothetical minimum. All additional sets are forgone choices so their payoffs are weighted by δ.

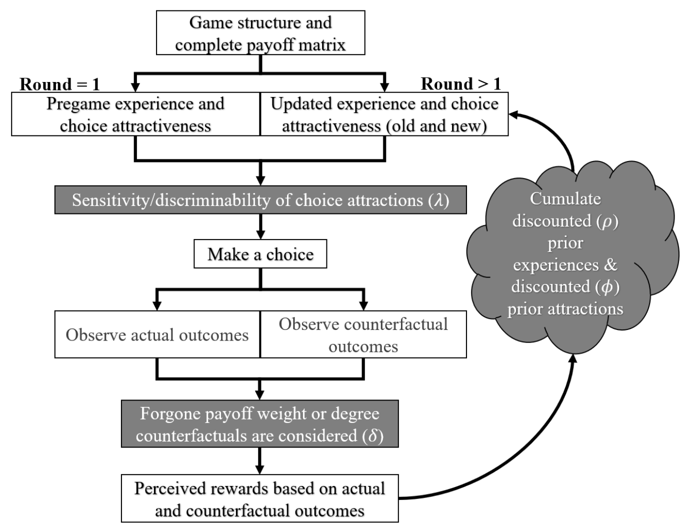

Figure 1: High-level summary of EWA behavior in the MEG.

EWA and EEWA were previously fit to MEG data (Hough, 2021) and starting choice attractions were based on sampling a choice from the first round choice distribution of the human data. We chose to estimate the mean and standard deviation for each parameter to fit average effort and intra-group variance. This introduced variability (i.e., individual differences) as parameter values for each agent were sampled from a normal distribution using estimated values. EWA and EEWA used the same estimated mean and standard deviation (parentheses): $\delta = .2(.01)$, $\rho = .9(.01)$, $\phi = .21(.17)$, and $\lambda = .49(.1)$. After model fitting, we found several issues. Both had a large portion of agents that never varied choices (35% compared to 1.4% of humans) suggesting an artificial fit. This is likely related to the low value of δ (.2), which conflicts with Byrne (2016), who suggest counterfactuals are weighted less than actual outcomes, but have a strong influence on behavior. Similarly, Camerer and Ho (1999) estimated δ as .85 for a similar game. A value of .2 makes actual payoffs carry 5x more weight than forgone ones and is likely to encourage repeated choices. The models were also missing features from the literature (e.g, beliefs about other players, initial preferences, strategy use and switching). Lastly, EEWA had a better fit than EWA, but it considers 27 forgone choices which might be unrealistic. These limitations motivated the development of a cognitive model that included the missing features and a higher degree of psychological correspondence. The model, called prediction, strategy, and simulation (i.e., PSS) included: player types, player choice predictions, strategies, and counterfactual thinking.

The PSS Model

The PSS model was implemented in the ACT-R cognitive architecture (Anderson, 2007), which includes both symbolic and sub-symbolic structures, and modules that represent systems of the mind. The PSS model uses the goal, imaginal, declarative and procedural modules. The imaginal module represents visual short-term or working memory and the goal module controls the model's current focus. The declarative memory module represents facts stored as chunks in long term memory and the availability of chunks is controlled by a sub-symbolic component. The procedural module uses condition-action rules (i.e., productions) to represent knowledge about how to do things. The procedural module's pattern matcher determines whether any production conditions match the current state and, if so, it "fires" and changes the state of the model. The PSS model includes three main features: predictions about other players, strategies, and learning. The instance-based learning framework (Gonzalez, Lerch, & Lebiere, 2003) is used for player choice predictions, however, we use a slightly different approach. Instances (i.e., chunks) are used strictly for player choice predictions, there is no pre-decision consideration of possible decisions, and decisions are a function of productions. Player choice predictions serve as input to strategies, which together produce a decision. After a decision is made, counterfactual thinking takes place and the model simulates unchosen strategies.

Players make choices simultaneously and often react to previous round choices. To capture reactions, instances store choices for two rounds: situation ($t-1$) and reaction choice sets (t). To leverage reactions, the model uses last round choices as cues to retrieve an instance(s) with best matching situation choices ($t-1$) (i.e., target), and the reaction choices (t) are extracted. The blending mechanism (Lebiere, 1999) aggregates accumulated reaction choices to serve as player choice predictions. Instances with a higher likelihood of retrieval, determined by activation, carry more weight. The ACT-R activation equation, $A_i = B_i + S_i + P_i + \varepsilon_i$, includes a: 1) base level term, B_i, for recency and frequency of use, 2) spreading term, S_i, for context effects, 3) partial matching term, P_i, for degree of match with retrieval cues, and 4) noise term, ε_i, for noise in memory. However, PSS uses blending instead of retrieval and only includes partial matching, P_i, and 4) noise, ε_i, terms. The blending mechanism uses an equation, $V = min \Sigma_i P_i * (1 - sim(V,V_i))^2$, to produce a value that minimizes the sum of all squared dissimilarities, $(1 - sim(V,V_i))^2$, of each chunk, i, and weights it by its probability of retrieval, $P_i = (e^{M_i/t})/(\Sigma_j e^{M_j/t})$. The probability of retrieval is a function of the match score for a chunk, $e^{M_i/t}$, which normalized by the match score of all retrieved chunks, $\Sigma_j e^{M_j/t}$. If the blended chunk has an activation below the activation threshold (default of 0), it fails and previous round choices serve as player choice predictions.

The PSS model includes four strategies (i.e., productions) that use player choice predictions to make decisions based on the 1) minimum, 2) average, 3) maximum, or 4) one higher than the average (i.e., signaling). After making a decision, the model receives feedback and updates the utility of strategies using payoffs as rewards in the ACT-R utility learning equation: $U_i(n) = U_i(n-1) + \alpha[R_i(n) - U_i(n-1)]$. The current utility for each strategy, $U_i(n)$, is a function of the: 1) previous utility, $U_i(n-1)$, 2) utility learning rate, α, 3) temporally discounted reward value, $R_i(n)$, and 4) a noise component, ε. Starting utility influences which strategies are initially selected, and the learning rate influences how quickly utilities change. The PSS model includes risk (i.e., RD) and payoff dominant (i.e., PD) player types (Van Huyck et al., 1990, 1991), which are represented by a pattern of starting utilities. RD players are risk averse and PD are willing to take risk in the pursuit of higher rewards. The four strategies were sorted by risk (i.e., min, ave, max, and signal) and RD players had the highest utility for the min-strategy (i.e., 130), which linearly decreased along the continuum to the signal-strategy (i.e., 70). The PD player type was defined as the opposite of RD. All unchosen strategies simulated during counterfactual thinking receive a fraction of the forgone payoff to correspond with psychology (Byrne, 2016) and game theory literature (Camerer, 2003).

PSS model parameter values were set based on ACT-R defaults, corresponding literature, or MEG structure. There were two architectural parameters for declarative memory: Activation noise and partial matching. Activation noise was set at its default value (i.e., 1). The mismatch penalty parameter (mp) was set to a small value (i.e., 1) so that all instances have influence. Procedural memory included two architectural (i.e., learning rate and noise) and two theory-driven parameters (i.e., starting utilities and counterfactual weight). Utility noise was scaled up from 1 to 7.5 to better correspond to payoff values and utility learning rate was left at the default value of .2. A counterfactual weight parameter, set to .75, discounted forgone payoffs (i.e., 75% of payoff) for strategies during counterfactual thinking. Two player types (i.e., RD and PD) had different starting utility patterns.

The model starts the game by selecting a player type, then predicts other player choices, makes a choice, and processes the results. For the first round, predictions and choices are sampled from the first round choice distribution of the human data. For all subsequent rounds (Figure 2), the model starts each round by attempting to recall and blend past instances with situation choices ($t-1$) similar to last round choices stored in the goal buffer. If successful, reaction choices (t) are blended and serve as player choice predictions. If blending fails, last round choices serve as player choice predictions. A new instance is then created in the imaginal buffer to store situation choices (i.e., last round choices), and predictions replace last round choices in the goal buffer. The model

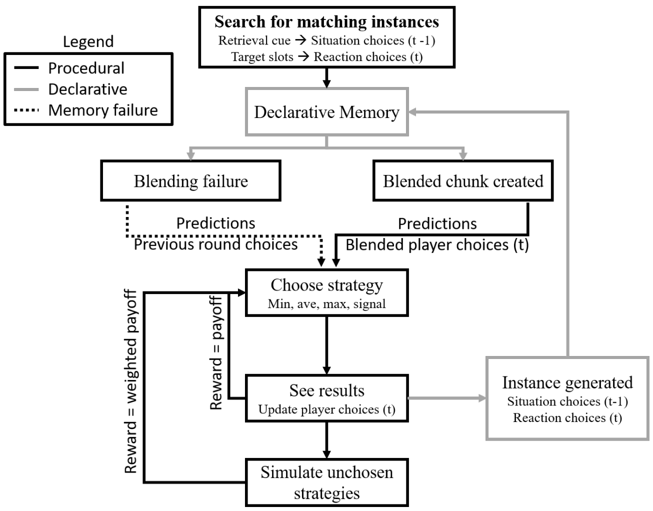

Figure 2: Simplified diagram of the PSS model processes.

selects the strategy with the highest utility, uses player choice predictions as input, and makes a choice. The model is then shown all player choices and its own payoff (i.e., results). At this point, a reward is triggered equal to the earned payoff and the utility of the chosen strategy is updated. In addition, values (i.e., reaction choices, decision, and payoff) are added to complete the instance in the imaginal buffer. The player choice predictions in the goal buffer are used to simulate all unchosen strategies to produce forgone choices (i.e., counterfactual thinking). Actual player choices in the imaginal buffer determine the forgone payoffs weighted by the counterfactual weight parameter. After unchosen strategies are simulated, the model replaces player choice predictions with actual choices in the goal buffer, and the instance is cleared from the imaginal buffer and is added to declarative memory.

Model Fit and Findings

EWA, EEWA, and PSS models simulated 100 groups with four agents playing the MEG and were fit to the average effort and intra-group variance of the human data (18 groups of 4).

The PSS model fit to average effort (Figure 3a), $r(38) = .4, RMSE = .27$, was better than EWA, $r(38) = .52, RMSE = .96$, and only slightly better than EEWA, $r(38) = .55, RMSE = .31$. However, EEWA had the best fit to average intra-group variance (Figure 3b), $r(38) = -.06, RMSE = .42$, followed by EWA, $r(38) = -.10, RMSE = .53$, and PSS, $r(38) = -.14, RMSE = 1.62$. EEWA also had the best fit to average payoff (Figure 3c), $r(38) = .43, RMSE = 7.19$, followed by PSS, $r(38) = .34, RMSE = 10.4$, and then EWA, $r(38) = .64, RMSE = 11.26$. EEWA best approximated average behavior, as PSS failed to fit intra-group variance. To better understand the data, we calculated variance and distance from the minimum (i.e., min-dist) for each individual. We found 35% of EWA and EEWA agents had choice variance of 0 and a mode of 0. About 1.4% of humans and 2% of PSS agents had no variance and modes were .27 and .05, respectively. Agent first round choices were based on first round choices of humans, with a mean of 4.38 and variance of 2.77. This lack of choice change likely contributed to EWA and EEWA's fit to both average effort and intra-group variance.

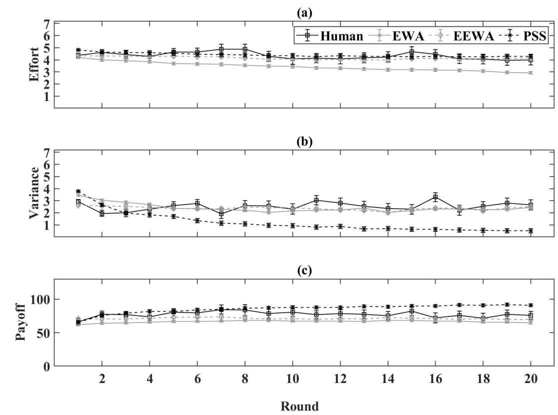

Figure 3: EWA, EEWA, and PSS model fits to the human data for average effort (a), intra-group variance (b), and payoff (c).

Next, we classified players as signalers if min-dist was above 0 for five consecutive rounds, as persistent signaling is more effective (Brandts et al., 2015). About 46% of humans were signalers, compared to 72% EWA, 78% EEWA, and 44% PSS agents. Most PSS agent signalers were PD types (64% compared to 17% for RD). Groups were then categorized based on amount of signalers and compared to assess how signaling influences group behavior. For simplicity, we classified groups as having 0-1 (i.e., -2), 2 , or 3-4 (i.e., 2+) signalers. Most human groups were classified as 2 (55.5%), followed by -2 (27.7%), and 2+ (16.6%). PSS had a similar pattern (44%, 34%, and 22%, respectively). For EWA and EEWA, most groups were 2+ (74% and 82%), followed by 2 (19% and 16%), and -2 (7% and 2%).

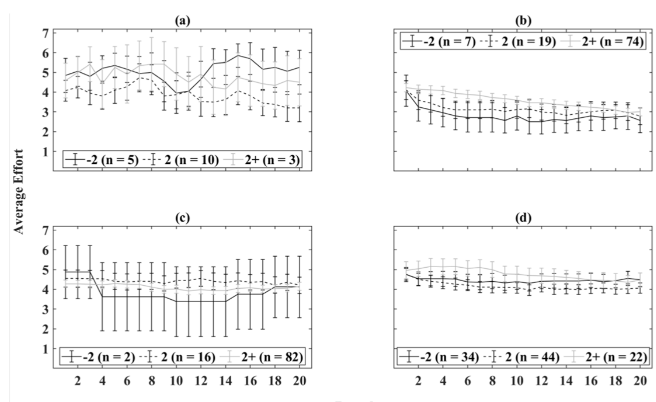

Figure 4: Average effort for signaler groups with human data (a), and EWA (b), EEWA (c), and PSS (d) model data.

We used linear mixed effects models with round and signaler group as fixed effects, and players nested within groups and groups as random effects. We report the most relevant interactions for comparison. For effort (Figure 4), human (a) -2 groups had the strongest positive trend and overtook the 2+ group ($Round * 2Group : \beta = -.07, t(1434) = -4.16, p < .001$ and $Round * 2 + Group : \beta = -.06, t(1434) = -2.58, p = .01$), EWA (b) 2+ groups had the strongest negative trend ($Round * 2 + Group : \beta = -.03, t(7994) = -3.63, p < .001$), and PSS (d) -2 groups had the weakest negative trend and overtook the 2+ group ($Round * 2Group : \beta = -.02, t(7994) = -4.98, p < .001$, and $Round * 2 + Group : \beta = -.04, t(7994) = -7.33, p < .001$). PSS had the most similar patterns for effort across signaler groups.

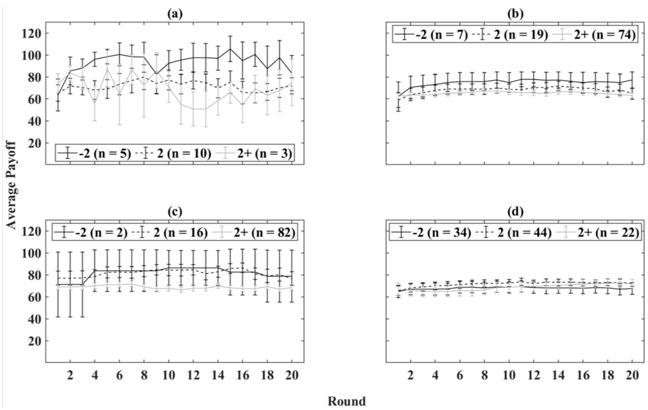

Figure 5: Average payoff for signaler groups with human data (a), and EWA (b), EEWA (c), and PSS (d) model data.

For payoff (Figure 5), human (a) -2 had the strongest positive trend ($Round * 2Group : \beta = -.56, t(1434) = -2.31, p = .02$ and $Round * 2 + Group : \beta = -1.33, t(1434) = -4.06, p < .001$), EWA (b) 2+ groups had the weakest positive trend ($Round * 2 + Group : \beta = -.32, t(7994) = -3.38, p < .001$), and PSS (d) -2 groups had the weakest postive trend ($Round * 2Group : \beta = .22, t(7994) = 4.98, p < .001$, and $Round * 2 + Group : \beta = .38, t(7994) = 7.33, p < .001$). For payoff, EWA and EEWA (no findings) were more visually similar to human data and PSS differed as payoff increased for groups with more signalers, suggesting effective signaling.

The signaler group results showed human groups with fewer signalers had higher effort and payoff, and Figure 6 shows this at group and individual levels. Coordination and response to signals was better with 0 (a) or 1 (b) signalers, and mixed with 2 (c) and 3 signalers (d). These groups suggest player behavior is complex and may involve changes in strategies. To explore underlying mechanisms of choice changes, we present an example group for EEWA and PSS.

In the EEWA group (Figure 7), one agent (P4) never varied its choice, two others made one change midway (P3) or at the last round (P1), and one made more than one change (P2). Effort and resulting low payoffs (b) show the "sticky choice" problem goes beyond no variance players. Choice attractions for agents 1 (c) and 2 (d) show the relationship between choice attractions and choice changes. One choice attraction dominates, suggesting the weighted choice sampling was necessary for choice changes.

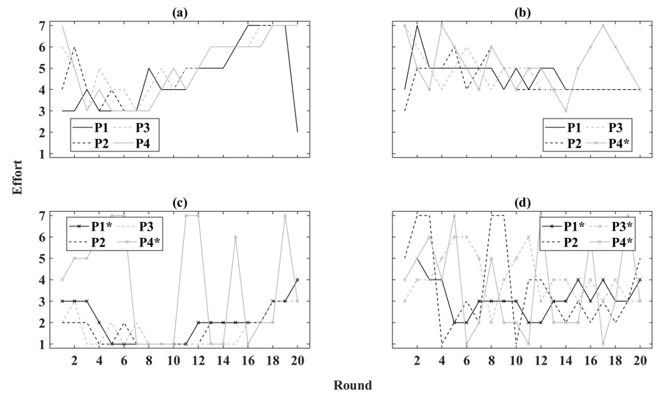

Figure 6: Player effort for human groups with 0 (a), 1 (b), 2 (c), and 3 (d) signalers.

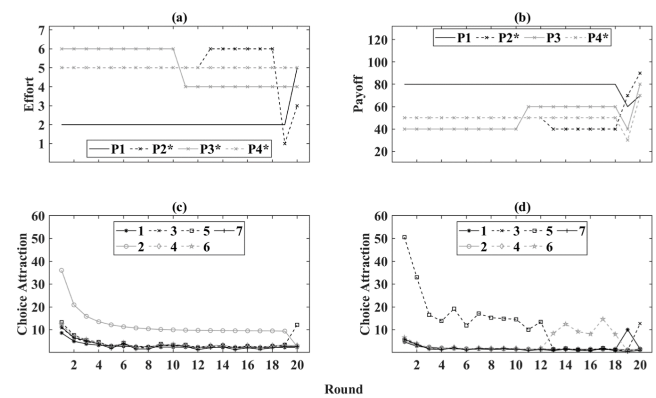

Figure 7: Agent effort (a) and payoff (b) for one EEWA group, with choice attractions for agents 1 (c) and 2 (d).

In the PSS group (Figure 8), there was more variation in choices and strategy competition. Agent 1 (P1) was a payoff dominant signaler that started at high effort, then ended on the minimum. Agent 3 (P3) was a risk dominant non-signaler that frequently set the minimum, then increased effort choices. Payoffs (b) show the benefits of these strategy shifts, most notably when agent 3 started choosing higher. Agent 1 (c) had the highest starting utility for signaling, then min- and ave-strategies competed until the min-strategy won. Agent 3 (d) had higher utility for the min-strategy, then competed with the ave-strategy. PSS Model agents demonstrated dynamic behavior by shifting from starting strategies based on group dynamics, counterfactual thinking, and learning.

Discussion

The complexities of coordination behavior and differences between models were only apparent through analyses at different levels of behavior. Results suggest analyses and modeling excluding individual and group level behavior(s) may

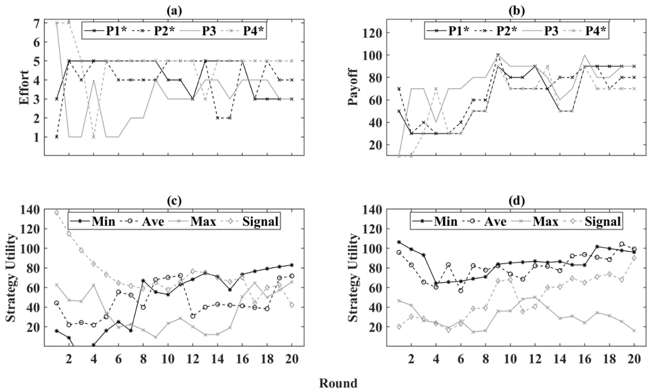

Figure 8: Agent effort (a) and payoff (b) for one PSS group, with strategy utilities for agents 1 (c) and 3 (d).

lead to faulty conclusions about behavior and underlying mechanisms. The PSS model took a step towards addressing these issues. It included features missing from previous models: player types, strategies, and player choice predictions. Although EEWA was the best fit to average behavior, PSS showed greater capability to approximate human behavior at average, group, and individual levels. PSS agents displayed interdependent behavior by switching strategies based on group behavior and learning, and model mechanisms allowed for explanation of each players behavior based on player choice predictions and strategy utility. However, there are several limitations. 1) Choice variation was approximated with arbitrary strategies, agents predicted their own choices to enable repeated choices, and signaling behavior was rigid. 2) The PSS model players were more sensitive to signaling and better able to coordinate, which corresponded with the literature, but not with the human data. 3) The player choice predictions were based on reaction choices and did not include forming beliefs about other's player types suggested in the literature. 4) Model comparisons were based on simulations and complexity was not punished, potentially giving PSS an unfair advantage. Overall, we highlighted issues related to levels of analysis, and the strengths and weaknesses of the PSS model can inform future work to better approximate and explain complex coordination behavior.

Acknowledgments

This research was supported by the U. S. Air Force Research Laboratory's 711th Human Performance Wing, Cognitive Models Branch. The contents have been reviewed and deemed Distribution A. Approved for public release. Case number: AFRL-2022-1525. The views expressed are those of the authors and do not reflect the official guidance or position of the United States Government, the Department of Defense, or the United States Air Force.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: Oxford University Press.

Bortolotti, S., Devetag, G., & Ortmann, A. (2016). Group incentives or individual incentives? a real-effort weak-link experiment. *Journal of Economic Psychology*, *56*, 60–73.

Brandts, J., Cooper, D. J., & Weber, R. A. (2015). Legitimacy, communication, and leadership in the turnaround game. *Management Science*, *61*(11), 2627–2645.

Byrne, R. M. (2016). Counterfactual thought. *Annual review of psychology*, *67*, 135–157.

Cachon, G. P., & Camerer, C. F. (1996). Loss-avoidance and forward induction in experimental coordination games. *The Quarterly Journal of Economics*, *111*(1), 165–194.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. New York: Russell Sage Foundation.

Camerer, C. F., & Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*(4), 827–874.

Costa-Gomes, M. A., Crawford, V. P., & Iriberri, N. (2009). Comparing models of strategic thinking in van huyck, battalio, and beil's coordination games. *Journal of the European Economic Association*, *7*(2-3), 365–376.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*(4), 591–635.

Hough, A. R. (2021). *Understanding and improving coordination efficiency in the minimum effort game: Counterfactual-and behavioral-based nudging and cognitive modeling*. Unpublished doctoral dissertation, Wright State University.

Hough, A. R., O'Neill, K., & Juvina, I. (2021). Counterfactual-based nudging and signaling promote more efficient coordination during group tasks. *Comprehensive Results in Social Psychology*, 1–27.

Lebiere, C. (1999). Blending: An act-r mechanism for aggregate retrievals. In *Proceedings of the sixth annual act-r workshop, george mason university, fairfax, va, usa.*

Leng, A., Friesen, L., Kalayci, K., & Man, P. (2018). A minimum effort coordination game experiment in continuous time. *Experimental Economics*, *21*(3), 549–572.

Luce, R. D. (1956). Semiorders and a theory of utility discrimination. *Econometrica, Journal of the Econometric Society*, 178–191.

Mehta, J., Starmer, C., & Sugden, R. (1994). The nature of salience: An experimental investigation of pure coordination games. *The American Economic Review*, *84*(3), 658–673.

Offerman, T. (2002). Hurting hurts more than helping helps. *European Economic Review*, *46*(8), 1423–1437.

Riechmann, T., & Weimann, J. (2008). Competition as a coordination device: Experimental evidence from a minimum effort coordination game. *European Journal of Political Economy*, *24*(2), 437–454.

Siegler, R. S. (1987). The perils of averaging data over strategies: An example from children's addition. *Journal of Experimental Psychology: General*, *116*(3), 250.

Van Huyck, J. B., Battalio, R. C., & Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, *80*(1), 234–248.

Van Huyck, J. B., Battalio, R. C., & Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, *106*(3), 885–910.