

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Understanding and Modeling Coordination in the Minimum Effort Game

Permalink

<https://escholarship.org/uc/item/60v1m2kw>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Hough, Alexander
Juvina, Ion

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Understanding and Modeling Coordination in the Minimum Effort Game

Alexander R. Hough (alexander.hough.1@us.af.mil)

Air Force Research Laboratory
Wright-Patterson AFB

Ion Juvina (ion.juvina@wright.edu)

Wright State University
Dayton, OH

Abstract

Groups of individuals need to coordinate in many real world domains. However, coordination failure is common and not well understood. There are few coordination measurements, analyses focus on averaged data, and models lack coordination strategies and clear correspondence to cognitive mechanisms. Here, we present a thorough analysis of human data from a difficult coordination scenario and a cognitive model implemented within the ACT-R cognitive architecture to fit and explain the data. Data were explored to better understand coordination strategies and group dynamics. The cognitive model included pre-game preferences, coordination strategies like signaling, and other player choice predictions. This work highlights the need for deeper data explorations and presents challenges for modeling related to coordination dynamics, strategies, and how players form beliefs about others.

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

Introduction

Humans often fail to coordinate in situations with competing individual and group incentives (Cooper, DeJong, Forsythe, & Ross, 1990, 1994; Camerer, 2003; Riechmann & Weimann, 2008; Van Huyck, Battalio, & Beil, 1990, 1991), which is often attributed to the lack of focal points or equally salient choices leading to optimal outcomes for all players (Blume, DeJong, Kim, & Sprinkle, 1998; Mehta, Starmer, & Sugden, 1994). This failure can result from miscoordination and/or inefficient coordination (Riechmann & Weimann, 2008). Coordination is the degree players settle or converge on a single choice and can be expressed numerically by calculating variance of choices within groups (Hough, 2021; Hough, O'Neill, & Juvina, 2021). Coordination efficiency involves where the outcome falls on a hypothetical continuum between worst and best possible outcomes.

Coordination often occurs over time, however, reaching and increasing coordination efficiency over time is more difficult (Brandts & Cooper, 2006; Brandts, Cooper, & Weber, 2015; Brandts, Cooper, Fatas, & Qi, 2016; Chaudhuri, Schotter, & Sopher, 2009; Van Huyck et al., 1991). Several techniques were applied to this issue, but involve changing the game structure (Chaudhuri et al., 2009; Van Huyck, Gillette, & Battalio, 1992; Van Huyck, Battalio, & Beil, 1993; Brandts et al., 2016, 2015). There are less invasive methods that are more generalizable. For instance, counterfactual thinking can increase coordination efficiency when it highlights outcomes that could have happened if choices were

more efficient (Hough et al., 2021). In addition, players can nudge each other to make more efficient choices by signaling or making choices that would result in better outcomes for everyone if the counterparts also made that choice. Signaling is more effective when it incurs a cost that others are aware of (Spence, 1978) and when it's persistent (Brandts et al., 2015, 2016). However, signaling is risky (Cachon & Camerer, 1996) and players often give up if it is not effective (Charness, Gneezy, & Henderson, 2018). Players also form and update beliefs about others to predict future behavior (Camerer, 2003), which can increase coordination.

To better understand coordination, coordination efficiency, strategies, and group dynamics, experimental data were analyzed and a cognitive model was developed to fit and explain the data. The experiment used a minimum effort game (MEG) (Van Huyck et al., 1990) to simulate a coordination scenario with: 1) simultaneous choices, 2) no communication, and 3) 20 rounds. Coordination is very challenging in this "weak link" game without an initial explicit focal point (Blume et al., 1998; Mehta et al., 1994) and other players choices, particularly the minimum, can become the focal point and influence players to coordinate on an inefficient choice (Brandts et al., 2015, 2016; Van Huyck et al., 1990).

The MEG

In the MEG, players make an effort choice between one and seven, and each player's payoff is determined by their choice and the group minimum (Table 1).

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

There are seven coordination points or Nash equilibria (Nash, 1951), which specify what a rational player should

select to maximize their own payoff regardless of counterpart choices (Camerer, 2003). They are represented diagonally from the payoff of 70 to 130. Nash equilibria are “Pareto ranked” according to payoff. A Pareto equilibrium maximizes the sum of payoffs for all players (Camerer, 2003). In terms of efficiency, Nash equilibria is efficient for the individual and Pareto for the group. In the literature and this paper, efficiency refers to group efficiency.

van Huyck et al. (1990, 1991) suggested players start by using risk or payoff dominant strategies, but deviate over time due to learning. Payoff dominance is a high risk, high reward strategy. The highest choice of seven can result in either the highest or the lowest payoff. Risk dominance is a low risk, low reward strategy. Choosing effort level one always results in the same payoff regardless of other player’s choices. These strategies serve as focal points, and in subsequent rounds, the minimum may become a salient focal point or anchor (Leng, Friesen, Kalayci, & Man, 2018; Van Huyck et al., 1990). This is a simple explanation for the frequently observed negative trend in effort over time across various manipulations like outcome information and group size (Camerer, 2003; Leng et al., 2018; Van Huyck et al., 1990, 1991).

MEG experiments (Bortolotti, Devetag, & Ortmann, 2016; Leng et al., 2018; Van Huyck et al., 1990, 1991) typically focus on effort and the minimum to analyze efficiency and signaling behavior. Leng et al. (2018) went a step further and identified signaling as alternating between the minimum and higher effort and found small increases in efficiency or the minimum. Bortolotti et al. (2016) identified weak links and found they were the source of coordination failure, but only early in the game. Despite these contributions, there is a lack of effective measurements, little is understood about coordination strategies and group dynamics, and no existing model is capable of capturing complex coordination behavior in the MEG. Specifically, player preferences, strategies like signaling, and beliefs about other players. In the following sections, a thorough data analysis and novel cognitive model are presented to better understand behavior in the MEG.

The MEG Experiment

A MEG experiment was conducted at a Midwestern University with 18 four-person groups (Hough, 2021; Hough et al., 2021). After all players made choices, they were shown all player choices and counterfactuals for both lower and higher choices and minimums. Players were not informed about the 20 round game length to reduce potential end effects.

Average effort and intra-group variance are plotted in Figure 1a. Averages were calculated per round for each group, then averaged across groups. Average effort is typically used to measure coordination efficiency and here, we use average intra-group variance to measure coordination (Hough, 2021; Hough et al., 2021). Notice variance is rather high and stable across 20 rounds, suggesting players did not coordinate well.

Average payoff (Figure 1b) carries information about average effort and the minimum. Average effort is stable around 4




Figure 1: Average effort and variance (a), and Average payoff (b). Error bars are 95% CIs.

and average payoff around 80. According to the payoff matrix (Table 1), this means the minimum was around 3. The minimum being one lower than the average suggests the presence of signaling. Calculating the distance from the minimum for individuals can indicate the presence and strength of signaling, identify group weak links, and categorize players as signalers if they choose higher than the minimum for 5 consecutive rounds (Hough, 2021; Hough et al., 2021). This signaler categorization is based on literature suggesting that persistent signaling is more effective (Brandts et al., 2015, 2016).




Figure 2: Average effort (a) and payoffs (b) for groups with -2, 2, and 2+ signalers. Error bars are 95% CIs.

Effective signaling should result in increases in effort and payoffs over time. This notion was explored by categorizing players as signalers (46% were signalers), then categorizing groups with less than 2 (-2), 2, or more than 2 sig-

nalers (2+). Here, a mixed effects model with random effects for players nested within groups and groups (Hough, 2021; Hough et al., 2021) was used to compare effort and payoff for group categories (Figure 2). For effort, there was an interaction effect with round for 2 ($\beta = -.07, t(1434) = -4.16, p < .001$) and 2+ ($\beta = -.06, t(1434) = -2.58, p = .01$) signaler groups, meaning -2 signaler groups had a more positive trend across rounds. For payoff, there was an effect for round ($\beta = .51, t(1434) = 2.53, p = .01$) and interaction effects with round for 2 ($\beta = -.56, t(1434) = -2.31, p = .02$) and 2+ ($\beta = -1.33, t(1434) = -4.06, p < .001$) signaler groups. This means there was a positive trend across rounds, which was more positive for -2 signaler groups.

The literature suggests persistent signaling should be effective. Here, we see groups with the fewest signalers have higher effort and payoffs. To better understand this, two groups are shown in Figure 3, a group without signalers (a) and a group with three signalers (b). The no signaler group (a) has better coordination, fewer signals, lower signaling magnitude, and increases in efficiency. The group with 3 signalers (b) is the opposite. Players appear to take turns signaling, setting the minimum, and even choosing lower than the previous minimum. Players differ in signaling frequency, magnitude, persistence, and response to other's signaling, suggesting they have different strategies. To better understand this complex behavior, a cognitive model was developed.




Figure 3: Average effort for players within a group with 0 (a) and 3 signalers (b) and all signalers marked (*).

A Cognitive Model of the MEG

The cognitive model includes: player types, player choice predictions, strategies, and counterfactual thinking, and agents that simulate humans playing the MEG. The model, referred to as the prediction, strategy, and simulation model (i.e., PSS), was developed in ACT-R. ACT-R is a hybrid cognitive architecture with both symbolic and sub-symbolic structures (Anderson & Lebiere, 1998; Anderson, 2007).

There are perceptual-motor modules (e.g., goal and imaginal) and two types of memory modules that represent systems of the mind. The PSS model uses the goal, imaginal, declarative, and procedural modules. The goal module determines the model's current focus and the imaginal module is used to temporarily store visual information. The declarative memory module represents facts in long term memory stored as chunks and a sub-symbolic component determines the availability of the chunks. The procedural module represents knowledge about how to do things, represented as condition-action rules. The pattern matcher of the procedural module determines which, if any, rules match the current state. If the condition of the rule matches the current state, then it "fires" and the action changes the state of the model. The behavior of a model is represented as a series of rule firings and corresponding changes to the state of the model.

The PSS Model

The main components of the PSS model include: predictions about other players, strategies, and learning. The instance-based learning approach (Gonzalez, Lerch, & Lebiere, 2003) is used as a framework for player choice predictions. However, the PSS model uses a slightly different approach (Juvina, Lebiere, & Gonzalez, 2015). Instances are used to make predictions about other players, decisions remain a function of procedural memory, and there is no pre-decision stopping rule for consideration of all possible decisions. A strategy is chosen, which uses player choice predictions to make a decision, then the unchosen strategies are simulated to represent counterfactual thinking.

Player Choice Predictions In the MEG, players make choices simultaneously and often have delayed reactions to previous round choices. Therefore, instances (i.e., chunks) contain player choices for two rounds: the previous ($t - 1$) and reaction choice sets (t). Last round choices are retrieval cues, previous choices ($t - 1$) are the target in the instance, and reaction choices (t) in the instance are the values of interest. Instances accumulate over time and the blending mechanism (Lebiere, 1999) was used to aggregate information to serve as player choice predictions. Every instance has an associated activation, that determines its likelihood of retrieval. Activation is determined by the activation equation: $A_i = B_i + S_i + P_i + \epsilon_i$. The activation of a chunk, A_i , is a function of the: 1) base level term, B_i , that represents recency and frequency of chunk use, 2) spreading term, S_i , that represents context effects, 3) partial matching term, P_i , that represents how well the chunk matches the retrieval cues, and 4) noise term, ϵ_i , that represents variability in memory. The PSS model uses blending instead of retrieval, so it only includes the partial matching, P_i , and noise, ϵ_i , terms. The blending mechanism retrieves a compromise value for all possible values of interest weighted by their probability of retrieval. The equation, $V = \min \sum_i P_i * (1 - \text{sim}(V, V_i))^2$, produces a value that minimizes the sum of all squared dissimilarities for values, $(1 - \text{sim}(V, V_i))^2$, of each chunk, i , and weights it by its

probability of retrieval, $P_i = (e^{M_{i/t}}) / (\sum_j e^{M_{j/t}})$. The probability of retrieval is a function of the match score for a chunk, $e^{M_{i/t}}$, that represents the degree of match between the retrieval cues and the target information in the chunk. The match score is normalized by the total match score of all retrieved chunks, $\sum_j e^{M_{j/t}}$. As player choices are integers, the compromise value is the sum of all chunk values of interest weighted by their probability of retrieval. Just like retrieval requests, blending can fail if the activation for the blended chunk is below the activation threshold (default of 0). If blending fails, the model uses previous round choices as player choice predictions.

Player Choice Strategies In order to make a choice, players use some kind of strategy or rule based on preferences or previous experience. The PSS model includes four strategies that use player choice predictions to determine choices: 1) the min-strategy selects the minimum, 2) the ave-strategy selects the average, 3) the max-strategy selects the highest, and 4) the signal-strategy selects one higher than the average. Note predictions included all players, so players included predictions of themselves. The model receives feedback and updates the utility of the chosen strategy based on the actual outcome. Strategy utility updates using the ACT-R utility learning equation: $U_i(n) = U_i(n-1) + \alpha[R_i(n) - U_i(n-1)]$. In the equation, utility for round n , $U_i(n)$, is a function of the: 1) previous utility, $U_i(n-1)$, 2) utility learning rate, α , and 3) current reward value, $R_i(n)$. There is an optional noise component, ϵ , that can be added to utilities, to add some stochasticity (Anderson, 2007). Starting utilities and the learning rate are important as they influence which strategies are initially selected and how quickly utilities change over time. The PSS model includes two patterns of starting utilities to represent risk and payoff dominant player types (Van Huyck et al., 1990, 1991). A risk dominant player (i.e., RD) is motivated to reduce risk or costs and would prefer choosing lower effort or stick to the minimum. A payoff dominant (i.e., PD) player is more willing to take risk and seek higher rewards. The RD player type was approximated by organizing strategies by risk (i.e., min, ave, max, and signal), setting the min-strategy at the highest payoff (i.e., 130), and linearly decreasing utilities along this continuum. The PD player type is defined as the opposite of risk dominant. In the PSS model, rewards are the payoffs accrued by strategies. After the chosen strategy receives a reward, the model engages in counterfactual thinking and simulates outcomes for the unchosen strategies in the same manner. However, they receive a fraction of the forgone payoff to better correspond to counterfactual thinking (Byrne, 2016; Kahneman & Miller, 1986) and the idea that it uses fictitious rather than actual outcomes (Camerer & Ho, 1999).

Model Overview There were two architectural parameters for declarative memory: partial matching and activation noise. Partial matching was a free parameter determined to be 1 through model fitting. There are no starting instances and there are only 20 rounds, so setting partial matching at 1 minimized mismatch penalties so that all chunks influence player

choice predictions. Activation noise is required for blending and was left at its default value of 1. For procedural memory, there were two fixed architectural parameters (i.e., learning rate and noise) and two parameters based on theoretically justified assumptions (i.e., starting utilities and counterfactual weight). The utility learning rate is set at the default value of .2 and utility noise was scaled up to 7.5 (i.e., default is 1) during model fitting to better correspond to payoff values (i.e., up to 130). The counterfactual weight (i.e., cfw) parameter was added to differentially weight payoffs for simulated strategies during counterfactual thinking. The cfw parameter was set to .75 so that counterfactual payoffs have 75% of the value as actual payoffs. There are two player types that correspond to a pattern of starting strategy utilities. In the final model run, there were more PD player types (57%), suggesting that there might be more PD than RD player types in the sample.




Figure 4: Simplified diagram of the PSS model processes.

During the first round, the model randomly selects a player type, predicts choices for the three other players, and makes a choice. Player choice predictions and the model's choice are randomly sampled from the first round choice distribution of the human data. The model skips situation recognition and instance blending (i.e., declarative memory), and strategy productions (i.e., procedural memory) in order to move directly to the results. For all subsequent rounds, the model goes through the same general process (Figure 4). The model attempts to recognize the situation by making a blending request using last round choices in the goal buffer as retrieval cues. The best matching chunks carry more weight in the blended choices and activation noise is added to the activation equation, which determines activation of the blended chunk. If blending is successful, the blended choices replace the last round choices in the goal buffer and represent player choice predictions. If blending is not successful, then last round choices serve as the player choice predictions. In both cases a chunk is created in the imaginal buffer to store last round choices in a new instance. Next, the model selects the strategy that has the highest utility and uses player choice predictions to make a choice. After all counterparts have

made a choice, the model moves to the results and is “shown” all player choices, its own payoff, and a reward is triggered equal to that payoff. The utility of the chosen strategy is updated using the utility equation: previous utility, the reward, the learning rate (i.e., .2), and the utility learning noise (i.e., 7.5). Player choices are also added to complete an instance chunk in the imaginal buffer. However, the predicted player choices remain in the goal buffer and the instance chunk remains in the imaginal buffer so that counterfactual thinking can take place. Next, the unchosen strategies are simulated one at time using the player choice predictions. The forgone payoffs are weighted by the *cfw* parameter (i.e., .75) and are used as the reward to update the utility of the strategy productions. Once all the unchosen strategies are simulated, the model stops counterfactual thinking. The model replaces the player choice predictions in the goal buffer with the actual choices from the current round for the next round. The instance chunk is then cleared from the imaginal buffer and is added to declarative memory. At this point, the round ends, and the model repeats the whole process for the next round.

Model Fit and Findings The PSS model was used to simulate 100 groups with four separate agents playing the MEG. It was fit to effort and variance of the human data. For comparison, we included a competing model from the game theory literature also fit to effort and variance, the Experience-weighted Attraction model (EWA) (Camerer & Ho, 1999).




Figure 5: PSS and EWA model fit to average effort and variance (a), and average payoff (b). Error bars are 95% CIs.

EWA is based on forming and updating attractions towards choices and features four parameters: 1) forgone payoff weight (δ) for all unmade choices in the same situation, 2) past attraction decay (ϕ) and 3) experience decay (ρ) that control the growth rate of choice attractions and recency effects, and 4) a discrimination sensitivity parameter (λ). EWA was previously fit to the human data and compared to the PSS model after the best fitting parameters were estimated (See

additional model comparisons in Hough (2021); Hough and Juvina (2022)). Figure 5 shows the model fit to average effort, variance, and payoff. The PSS model had a better fit to average effort, $r(38) = .4, RMSE = .27$, compared to the EWA model, $r(38) = .52, RMSE = .96$. However, EWA had a better fit to variance, $r(38) = -.10, RMSE = .53$, than the PSS model, $r(38) = -.14, RMSE = 1.62$. The PSS had a slightly better fit to payoff, $r(38) = .34, RMSE = 10.4$, than EWA, $r(38) = .64, RMSE = 11.26$. Further inspection revealed 35% of EWA players stuck to first round choices that were sampled from the first round choice distribution of the human data (like the PSS model). For comparison, only 1.4% of humans displayed this behavior and 2% of PSS players. This choice stickiness helps explain why EWA was better able to fit variance. Although the PSS model could not fit variance as well, its amount of choice stickiness was more similar to humans and it provided a better fit to effort and payoff. Note that the PSS model was implemented in ACT-R making it inherently more complex and it was not penalized for that here.




Figure 6: Average effort (a) and payoffs (b) for PSS model groups with -2, 2, and 2+ signalers. Error bars are 95% CIs.

Next, we look at group dynamics and relate them to strategy use. Model players were classified as signalers using the same technique as the human data. About 44% of model players were classified as signalers (compared to 46% of humans) and most of them were payoff dominant player types (64% compared to 17% for risk dominant). Next, we ran the same linear mixed effects models for effort and payoff for comparison (Figure 6). Similar to the human data, we found interaction effects for round and both 2, $\beta = -.02, t(7994) = -4.98, p < .001$, and 2+ signaler groups, $\beta = -.04, t(7994) = -7.33, p < .001$, meaning both 2 and 2+ signaler groups had a greater negative trend than -2 groups. For payoff, there were interaction effects for round and both 2, $\beta = .22, t(7994) = 4.98, p < .001$, and 2+ signaler groups, $\beta = .38, t(7994) = 7.33, p < .001$, meaning 2 and 2+ signaler groups had a more positive trend than -2 signaler groups. The

PSS model payoff results contrast with the human data where 2 and 2+ groups had a more negative trend than -2 groups. This suggests signaling was more effective for PSS model groups. Next, we explore PSS model group dynamics for one example group with 3 signalers (Figure 7).




Figure 7: Average effort and variance (a), and average payoff (b) for a PSS model group with 3 signalers marked (*).

Player 1 was a payoff dominant signaler and consistently choose higher effort up to round 9, then started choosing lower and ended the game choosing the minimum. Player 3 was a risk dominant non-signaler and set the minimum for the first 12 rounds, then gradually increased effort choices. Payoffs in Figure 7a show the benefits of these strategy shifts. Player 3 increased payoffs by setting a higher minimum, which increased payoffs for all other players. To better understand strategy shifts, Figure 8 shows strategy utility for players 1 (a) and 3 (b). Here, we can see the influence of counterfactual thinking and learning on strategy utilities over time. Player 1 had the highest starting utility for signaling, then the min- and ave-strategies started to compete, until the min-strategy eventually dominated. On the other hand, player 3 started out with higher utility for the min-strategy, then the min- and ave-strategy competed for the rest of the game. Model players demonstrated dynamic behavior by shifting from starting strategies based on group dynamics, counterfactual thinking, and learning. Furthermore, we can peer into the model to quantitatively understand and explain this behavior.

Discussion

We analyzed human MEG data under the assumptions that: 1) intra-group variance indicates coordination, 2) payoffs indicate coordination efficiency, and 3) distance from the minimum indicates signaling frequency and magnitude. We found suggestive evidence that participants classified as signalers used sophisticated strategies, beyond payoff and risk dominance, which presents a potentially valuable addition to the literature regarding signaling. However, it is still not clear




Figure 8: Strategy utilities for player 1 (a) and player 3 (b) from Figure 7.

how to appropriately measure and classify signaling behavior. The PSS model, implemented in ACT-R, was developed to better understand and track behavior. The PSS model included: player types, strategies, player choice predictions, and counterfactual thinking. The model better fit average effort and payoff compared to a competing model, but could not appropriately fit variance. PSS model players displayed dynamic and interdependent behavior by switching strategies over time based on group behavior and learning. The nature of the model allowed for explanation of each players behavior based on player choice predictions and changes in strategy utility. However, there are several limitations. 1) Choice variation was approximated with four arbitrary strategies. One strategy represented choosing the minimum predicted choice and three strategies represented different types of signaling behavior based on magnitude. There was no strategy for making the same choice as last round (e.g., setting the minimum again), which was handled by having players predict their own choices. Future work could improve signaling behavior, perhaps with a cost/reward function to control signaling and its magnitude, and players being able to make repeated choices. 2) The PSS model players were on average, more sensitive to signaling and better able to coordinate. It's not clear why coordination was so poor in the human data. More data would be helpful to better inform how to modify the model. 3) The player choice predictions were based on reaction choices from one round to the next. The literature often refers to players forming beliefs about other's player types based on patterns of behavior, which could be added. 4) Further PSS model work is necessary to determine which features are useful, how well it explains human behavior, and re-assess its fit relative to its additional complexity over EWA. Overall, this model served as a first step towards explaining data in the MEG beyond average statistics and its strengths and weaknesses can inform future modeling work.

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-2022-0394. 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. Human data collection and model development was part of Alex Hough's Dissertation at Wright State University. Participant compensation was supported by Wright State University, specifically, a Graduate Student Assembly grant awarded to Alex Hough and research funds from the Psychology Department. We thank Christopher Stevens and four anonymous reviewers for constructive feedback on an earlier draft of this paper.

References

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York, NY: Oxford University Press.
- Anderson, J. R., & Lebiere, C. J. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Blume, A., DeJong, D. V., Kim, Y.-G., & Sprinkle, G. B. (1998). Experimental evidence on the evolution of meaning of messages in sender-receiver games. *The American Economic Review*, 88(5), 1323–1340.
- 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. (2006). A change would do you good.... an experimental study on how to overcome coordination failure in organizations. *American Economic Review*, 96(3), 669–693.
- Brandts, J., Cooper, D. J., Fatas, E., & Qi, S. (2016). Stand by me—experiments on help and commitment in coordination games. *Management Science*, 62(10), 2916–2936.
- 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.
- Charness, G., Gneezy, U., & Henderson, A. (2018). Experimental methods: Measuring effort in economics experiments. *Journal of Economic Behavior & Organization*, 149, 74–87.
- Chaudhuri, A., Schotter, A., & Sopher, B. (2009). Talking ourselves to efficiency: Coordination in inter-generational minimum effort games with private, almost common and common knowledge of advice. *The Economic Journal*, 119(534), 91–122.
- Cooper, R. W., DeJong, D. V., Forsythe, R., & Ross, T. W. (1990). Selection criteria in coordination games: Some experimental results. *The American Economic Review*, 80(1), 218–233.
- Cooper, R. W., DeJong, D. V., Forsythe, R., & Ross, T. W. (1994). Alternative institutions for resolving coordination problems: Experimental evidence on forward induction and preplay communication. In *Problems of coordination in economic activity* (pp. 129–146). Springer.
- 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., & Juvina, I. (2022). Individual differences and levels of analysis in computational models of coordination [manuscript submitted for publication].
- 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.
- Juvina, I., Lebiere, C., & Gonzalez, C. (2015). Modeling trust dynamics in strategic interaction. *Journal of applied research in memory and cognition*, 4(3), 197–211.
- Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological review*, 93(2), 136.
- 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.
- 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.
- Nash, J. (1951). Non-cooperative games. *Annals of mathematics*, 286–295.
- 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.
- Spence, M. (1978). Job market signaling. In *Uncertainty in economics* (pp. 281–306). Elsevier.

- 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.
- Van Huyck, J. B., Battalio, R. C., & Beil, R. O. (1993). Asset markets as an equilibrium selection mechanism: Coordination failure, game form auctions, and tacit communication. *Games and Economic Behavior*, 5(3), 485–504.
- Van Huyck, J. B., Gillette, A. B., & Battalio, R. C. (1992). Credible assignments in coordination games. *Games and Economic Behavior*, 4(4), 606–626.