Going with the group in a competitive game of iterated reasoning

Seth Frey (sethfrey@indiana.edu)
Psychology Building, 1101 E. 10th St.
Indiana University, Bloomington, IN 47405

Robert L. Goldstone (rgoldsto@indiana.edu)
Psychology Building, 1101 E. 10th St.
Indiana University, Bloomington, IN 47405

Abstract

In some strategic games, thinking ahead about other players' reasoning can lead to better predictions about what they will do. In other games, infinitely iterated reasoning ultimately prescribes random play. In an online experiment of strategic thinking in groups, we tested participants in a game with the formal structure of a random game, but the superficial structure of a game that rewards iterated reasoning. We found that participants conformed to the superficial structure of the game, and earned more than they would have by playing randomly.

We estimated how many steps participants thought ahead in the game and discovered implicit coordination at the group level. Participants unexpectedly "matched" their degree of iterated thinking to each other.

Keywords: experiment; strategic thinking; thinking steps; collective behavior; beauty pageant; matching pennies; rock paper scissors; n-player games

Introduction

Reflecting on how people learn to think ahead, behavioral economist Colin Camerer muses that

> ...strategic thinking seems to be more like learning to windsurf, ski, or fly an airplane, activities that require people to learn skills which are unnatural but teachable, and less like weight-lifting or dunking a basketball, where performance is constrained by physical limits. (Camerer, 2003, p.[249])

Our work challenges this analogy to individual activity. We propose that thinking strategically—thinking through the anticipated actions of others—may be a group activity: less like skiing or dunking and more like hockey or basketball. We report an experiment in which participants set their level of strategic thinking to match that of their peers, and earned more than economic theory predicts.

The Nash equilibrium is a cornerstone of game theory. It describe those outcomes of an economic game for which no player can gain by changing their strategy individually. John Nash proved that every game has at least one such equilibrium.

Mapped onto the real outcomes it models, the reasoning process behind the Nash equilibrium invokes an infinite regress of thoughts about thoughts. But human limits to this ideal reasoning process constrain the applicability of theory to real decisions. This concern has made experiments about iterated reasoning important.

Ho, Camerer, and Weigelt (1998) have shown demographic and motivational effects on the number of steps that participants think ahead through the thoughts of others' thoughts. In unpublished data, they report that participants playing for higher stakes ($28 vs. $7) thought further ahead by about half of a thinking-step. Other work has shown that training can also improve iterated reasoning. In the project that motivated the introductory analogy to individual sports, Costa-Gomes, Crawford, and Broseta (2001) elicited near-perfect equilibrium play from participants who had been exposed to the relevant aspects of game theory before the experiment.

Matching Pennies

The novel game that we tested in this experiment has some neighbors among familiar economic games. The nearest neighbors are Matching Pennies—a game that rewards randomness—and the Beauty Pageant, a classic domain for studying the levels of iterated reasoning.

In Matching Pennies, two players select between two strategies: *heads* and *tails*. One player earns a point if the two selected strategies are the same (e.g. both *heads*), and the other player earns the point if they are different (e.g. one *heads* and one *tails*).

Matching Pennies can be considered a two-player version of the children's game Rock Paper Scissors. In these games there is no single best strategy: every choice can be defeated by a choice that can also be defeated. Rock Paper Scissors and Matching Pennies are the archetypal examples of games with mixed-strategy equilibria. The best approach in each of them is to pick strategies randomly (rather than playing "pure" strategies). Mixed strategies draw randomly from the possible strategies to maximize expected payoff. Behavioral results from the simplest two- and three-person Matching Pennies games are consistent with the hypothesis that people play mixed strategies (Goeree & Holt, 2001; McCabe, Mukherji, & Runkle, 2000).

The Beauty Pageant

Though the game that we report is formally related to Matching Pennies, it has a superficial resemblance to another popular economic game, the Beauty Pageant.

In the Beauty Pageant game, competing players are told to select a number 1–100. A player wins if they guessed the value that is closest to two-thirds the average of all other submitted values.

Reasoning that (a) no player would choose a value over sixty-six, (b) all players would realize this, (c) therefore, no player would chose a value over $2/3 * 66$, (d) all players would recognize *this*, and (e) so on, the player who iteratively applies strict dominance to the set of strategies will act consistently with the Nash equilibrium and select zero. Compare this with Matching Pennies, in which iterated reasoning leads to uniformly random play.

The Beauty Pageant can be used to study iterated reasoning, because if participants choose a value higher than zero, an investigator can infer the number of thinking steps that must have been behind that choice. A player thinking zero steps ahead will expect others to play randomly and select a value that is $2/3$ of this mean: 33. A player who selected $2/3$ of 33 must have been thinking one step ahead, and a player who was thinking two steps ahead will choose a value that is two-thirds of that.

In the first experimental studies of the Beauty Pageant, guesses were well above zero (Nagel, 1995). Nagel's results were consistent with the hypothesis that most participants think only one or two steps ahead. In subsequent rounds of play, participants all started guessing lower, and selections crept slowly down towards the equilibrium at zero.

In the first round of play, a persistent fraction of participants make choices that are larger than 33 and even larger than 50. By contrast, very few participants select the equilibrium response at zero. Pursuing this result, experimentalists are finding that even when participants recognize dominated strategies themselves, they tend to doubt that others will (Camerer, 2003). This doubt about the rationality of others may be what truncates the number of iterations that the average participant makes. Building from this work, we tested participants in a game with features of both Matching Pennies and the Beauty Pageant.

Method

Participants

We collected data from 154 psychology undergraduates participating in twenty-nine experimental sessions at Indiana University. Groups ranging from two to nine participants were split into cubicles to play a computerized game of strategic reasoning. We did not control group size. Most group sizes were represented in at least four sessions. The exceptions were at groups of size six, eight, and five, which we tested in one, two, and three sessions, respectively. Group size did not change significantly over the semester of data collection. Participants earned course credit for participating, and were motivated within the experiment to earn points. Points have been shown to be sufficient to elicit motivated behavior in experiments of economic behavior (Camerer & Hogarth, 1999). Post-experiment interviews were consistent with the idea that subjects enjoyed earning points.

Task

Participants played sixty rounds of the *Wheel Game*. In each round of play, they chose synchronously and blindly among twelve possible options, or *strategies*. The twelve strategies were displayed as a grid of numbered tiles (Figure 1). Participants earned one point when their selected strategy was immediately above the strategy of another participant.

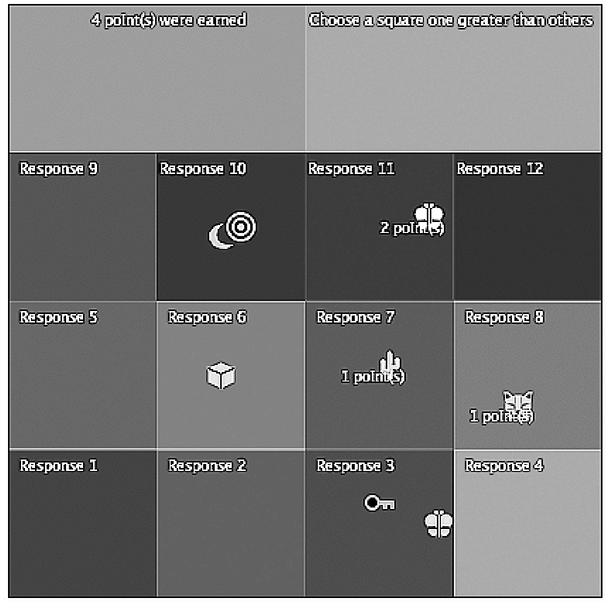

Figure 1. Screenshot of game board. Players earn one point for each competing player whose tile, or *strategy*, was one less than their own. The strategies wrap around such that the first is one-above the twelfth, and no strategy dominates any other. This picture was taken between rounds, when the positions of other players are revealed and points are distributed.

For example, if two players selected Tile 4, and three picked Tile 5, then each of the three players on the higher tile would earn two points. Subjects had no explicit incentive to avoid being one tile below another participant, only to be one tile above the others.

The game is called the Wheel Game because the set of strategies "wraps around" such that a participant on Tile 1 earns a point for each participant on Tile 12.

Procedure

All participants' decisions were revealed after each round by showing each participant's icon on the selected strategy tile. We did not reveal the choices of other groupmates until after all group members had selected a strategy tile. A box next to the game board displayed each participant's accumulated points as the rounds progressed.

The Wheel Game was one of a series of six games that groups played in each experiment. These other games will be reported elsewhere. Each participant's uniquely identifying icon persisted for each round of all six games. In each session we randomized the order of the six games. This controlled for the effects of learning and reputation across games, and not within games.

Predictions

In the simplest version of Matching Pennies, the unique Nash equilibrium is mixed: Each player selects both strategies with 50% probability. Scaling up to its twelve-stratecy analogue, the symmetric mixed-strategy equilibrium in the Wheel Game is for all players to draw uniformly from the twelve strategies. We assumed that randomly moving players would randomize uniformly, and this is the only mixed strategy that we tested for.

However, the Wheel Game is different from a twelve-player Matching Pennies game because the number of players in the groups we tested did not equal the number of strategies. Group sizes in the Wheel Games we tested were always smaller than the number of strategies. Because of this, only a fraction of available strategies were ever picked in a given round. *We suspected that participants would treat the Wheel Game less like a game of mixed strategies and more like a game of iterated reasoning*, like the Beauty Pageant.

The Wheel Game differs from the Beauty Pageant in several ways. First, there are only twelve strategies, and they wrap around. This precludes the existence of a unique pure strategy Nash equilibrium. Players earn the most in a round if they are above the tile with the most participants. They are thus attempting to predict the modal tile, and select the tile above it, with the knowledge that everyone else is attempting to do the same thing.

Earnings in the Wheel Game are not winner-take-all, they can be distributed among players in different ways. Since multiple players can earn points in a round, there is room for collusion between them—a possibility explored below.

Insofar as participants actually treat the Wheel Game like the Beauty Pageant, it allows similar inferences from participants' choices to their levels of iterated reasoning. If I was just scored on at Tile 4 by a single static peer at Tile 5, my next move might be to Tile 6. Expecting this reasoning of me, a peer might move to Tile 7. One step of iterated reasoning, on my part, would lead me to Tile 8, and two steps would take me up to Tile 10.

Unfortunately for this interpretation, a move from Tile 4 to Tile 8 might also result from looking ahead seven steps—all the way around the wheel. However, participants in many experimental settings have been shown to think ahead one to two steps on average (Camerer, Ho, & Chong, 2004, 2002). By the reasoning above, this average is not enough to reason entirely "around" the wheel of strategies. We suspected that participants would treat this game (or expect groupmates to treat this game) as one that evokes the "creeping" of strategies in iterations of experimental Beauty Pageants, rather than the random play observed during Matching Pennies.

Although we predicted that participants in this mixed-strategy game would behave as if they were in a different kind of game, like the dominance-solvable Beauty Pageant, we did not suspect that this behavior would collectively lead to higher scores, or that players in a group would start to match each others' thinking steps around the wheel.

Results

Ruling out the Symmetric Mixed Strategy Equilibrium

Our results show that behavior in the Wheel Game is not consistent with game-theoretic prescriptions for uniformly random play, or with previously observed behavior in two- and three-player Matching Pennies. The Nash equilibrium predicts that players will select randomly from among the twelve strategies. We did not observe this. As demonstrated below, the entropy (randomness) of participants' behavior was far below that of simulated random play, and we found evidence consistent with the hypothesis that participants tried to predict the choices of their peers—and that they collectively "rotated" around the strategy space. We compared the behavior of actual groups to that of simulated groups selecting uniformly from the twelve strategies at each round.

Our analysis revealed structure within the groups that rules out random play. First, we measured the randomness of individual behavior directly using Shannon's information entropy (Shannon & Weaver, 1949). The expected entropy, H, of sixty rounds of random strategy selection was 3.58 bits. Participants' behavior generated a lower entropy of 3 (Wilcoxon signed ranks $n = 154, V = 0, p < 0.001$).

Second, the visual layout of the game played some role in making participants' selections non-random. The most commonly selected tiles were Tile 12, Tile 1, and Tile 7. Third, a participants' choices showed a clear dependence on previous round, as elaborated below.

If participants are not playing with the randomness that the most salient mixed-strategy prescribes, it could be the case that they do not care about the game or understand it. However, participants scored points on each other significantly more often than would be expected by random players ($t(275.6) = 4.12, p < 0.001$). Their scoring rate was 0.576 points per round, compared to the 0.454 expected of random players. Randomness was not correlated with score. This makes it unlikely that people are bad randomizers who are learning the mixed strategy, or attempting to play it. A handful of very low entropy participants were scored on many times, but this was not common enough to register a significant trend in either direction.

Rotation

After ruling out random behavior consistent with the symmetric mixed equilibrium, we pursued our suspicion that participants would slowly rotate forward around the circle of strategies.

We were able to show evidence for consistent rotation at some rate, probably around 4.5 steps forward. To show this, we defined a participant's *rate* in a round as the change in their strategy from the last round. Because a participant's intent is not observable, we constrained rate to an integer from 0 to 11, which could represent an intended movement of $rate \times$ some integer i. For example, if a participant chose the same strategy in consecutive rounds, we interpreted this as a rate of zero, even if the subject imagined it as a change of twelve or some multiple of twelve.

49% of rate values (fifty-nine per subject per session) were between 0 and 5.5 (exclusive). 34% of rates were between 6.5 and 11. 17% of decisions were either zero or six steps away from the previous choice, and did not suggest rotation in either direction.

Though mean rates were consistently below 5.5, all groups showed spurts in which the mean rate exceeded 6 tiles per round. One group diverged down from that average rate of four tiles per round towards an average rate of one. Another group alternated wildly between rates of three and seven steps for much of the sixty rounds.

If participants had been playing the symmetric mixed Nash equilibrium, mean rate would have registered a value of 5.5 (half of eleven, the maximum rate). Observed mean rate was significantly lower, at 4.31 (Wilcoxon signed ranks $n = 308, V = 2780, p < 0.001$).

A Group Effect

We found that interactions between participants influenced their reasoning. Most interestingly, we found that participants' rates were correlated with the mean rate of the rest of their group ($t(5,303) = 5.34, p < 0.001$) (Figure 2). This implies that groups converged on a "group rate," and that each participant's own rate must have been some function of the observed behavior of the other participants in the group.

We also observed an increase in the rate of rotation by block (the first thirty and second thirty rounds of play). Rotation increased by half a tile, from 4.1 to 4.53 tiles per round ($F(8,299) = 12, p < 0.001$). This is consistent with the learning observed in popular games of iterated dominance, like the Centipede Game (Rapoport, Stein, Parco, & Nicholas, 2003) and the Beauty Pageant (Ho et al., 1998; Ho, Camerer, & Weigelt, 2003). The same researchers also found an increase in thinking steps with group size (Ho et al., 1998; Rapoport et al., 2003). We did not find a corresponding result.

Clustering

We observed other patterns in the structure of our groups that support the hypothesis of itereated reasoning. We reasoned that if players were trying to anticipate each others' moves around the circle, their choices would be clustered around each other. We used a simple measure of clustering for each round: the sum, over all players in a group, of the number of other players who have selected the same strategy in that round. By this measure, the following three strings of seven digits are progressively more clustered: $[1,1,1,1,1,1,1]$, $[0,3,3,1,0,0,0]$, $[0,0,7,0,0,0,0]$. Their clustering values are 0, 12, and 42, respectively. Observed clustering was significantly higher than in the random benchmark simulations (Wilcoxon signed ranks $n = 8, V = 36, p = 0.008$), but it decreased towards random across blocks (Wilcoxon signed ranks $n = 154, V = 9760, p < 0.001$). This change in clustering with time was not accompanied by a significant change in average performance. Other tests of clustering supported these results.

Collusion

For all group sizes that we tested, the maximum number of points that can be earned in round of the Wheel Game is $floor(n/2) * ceiling(n/2)$. This corresponds to a group strategy in which half of the group is one step above another half of the group. Since the Wheel Game is not winner-take-all, participants may have found a way to coordinate and earn more points. Of the coordinative strategies, some would permit individual payoffs that are both large and equitable. For example, if players had split into two groups they could have "leap-frogged" through the strategy circle. If participants had converged on this sophisticated strategy, then both their scores and their rates of being scored on would have been much higher.

We found a positive correlation between scoring and being scored on, controlling for group size ($t(9,299) = 2.12, p = 0.034$). However, we also found that mean earnings were only 32.5% of the maximum possible, and the highest percentage—from a group of size seven—was only 51.1% of this maximum.

Figure 2. Each point represents a participant. The axes plot each participant's mean rate against the mean of all other group members combined. The positive correlation implies that participants learned to match their rates of rotation—and thus levels of iterated reasoning—to each other.

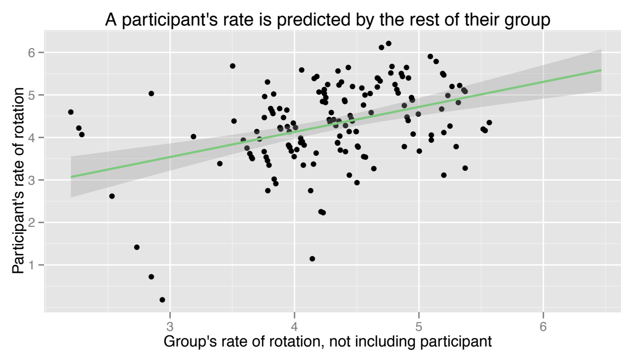

Discussion

Rotation

Although the Wheel Game is formally a descendant of Matching Pennies, groups behaved as if they were playing an infinitely receding Beauty Pageant—they rotated around the wheel of strategies instead of selecting uniformly from the twelve strategies.

We suggest that the average participant was rotating four-and-a-half steps forward, and thinking between one and two steps ahead. This behavior is not directly demonstrable in the Wheel Game. If a player moves from Tile 3 to Tile 6, there is no way to know if this was a step forward by three, backwards by nine, or even a leap of thirty-nine. However, the present data, and extant empirical and theoretical work all support this proposal. Participants expected each other's selected strategies to depend on previously selected strategies.

In their argument for the Poisson Cognitive Hierarchy model of strategic thinking, Camerer, Ho, and Chong (2004) review many past games and find evidence for a "universal inconstant" of 1.5 thinking steps. This implies that Joe will usually expect Sue to act strategically, but that Joe doesn't tend to think Sue will expect him to act strategically. Though many games elicit many different levels of iterated reasoning, most games elicit one to two levels of iteration, and the mean of 1.5 fits the data for many games impressively well. Parameterized at 1.5, and tested against data on a diverse set of over one hundred experimental games, Camerer et al.'s model outperforms every other general behavioral model.

If participants in the Wheel Game are anticipating another's guess of their own future move, then a rate between four and six tiles is consistent with the 1.5 thinking steps offered by Camerer et al.. This makes it more likely that participants are rotating forward at approximately four tiles per round, rather than forwards or backwards by four-plus- some multiple of twelve.

Past experiments have observed an increase in thinking steps with learning. If the rotation that we observe in the Wheel Game is due to the iterated reasoning of participants, increased experience in the game would lead to an increase in the average rate of rotation. We observed this effect, and the result supports our proposal that the same type of reasoning underlies the Wheel Game and more conventional experiments of iterated dominance.

It is possible that the Wheel Game elicits strategic thinking at higher levels than the average reported by Camerer and Ho. After all, participants in the game have an easy way to "look" smart. To think farther ahead, a participant only has to increment the strategy they are considering. Higher levels of iterated reasoning are not unheard of, Camerer (2003) reports unpublished evidence of participants thinking ahead as many as four and five steps ahead p.[18].

But even if participants are thinking entirely around the wheel, our conclusions are still valid. We do not even require that different participants are advancing by the same multiple of twelve. Our conclusions require only that (a) participants are strategizing about the strategizing of others to some depth, (b) thinking depth is correlated with the number of tiles a participant moves ahead each round and (c) participants are moving forwards rather than backwards around the wheel. These conditions are sufficient to support our claim that groups rotate around the strategies of the Wheel Game.

Together, our results demonstrate that behavior in the Wheel Game is inconsistent with the focal symmetric mixed Nash equilibrium. It is common for experimentalists to see this kind of result and triumphantly claim that they have "disproven" game theory. Other researchers are more cautious. Camerer reconciles the inconsistencies between theory and experiment by observing that, even when people miss the equilibrium, they trend towards it eventually. He makes the bold, carefully hedged, and carefully researched claim that "there are no games so complicated that participants do not converge *in the direction of* equilibrium (perhaps quite close to it) with enough experience in the lab." (emphasis added; Camerer, 2003p.[20])

Even this careful compromise may not be sufficient to account for our results. Behavior after sixty rounds of the Wheel Game is not only inconsistent with equilibrium, it has moved away from it. Mean entropy is much lower than what would be expected from mixed-strategy play, and it is decreasing further with time. Mean scores are higher than random. Participants seem to be converging upon some other behavioral regularity. If it turns out that the Wheel Game is not too complex, and that sixty rounds of play provide sufficient experience, then it constitutes an exception to Camerer's universal claim that people eventually behave more consistently with theory.

Despite the high scores and low randomness in strategy selection, there was some evidence supporting the idea that groups were approaching random play in the longer term. Clustering moved significantly in the direction of equilibrium with time. Scores decreased as well, though not significantly. Also, entropy at the group level (entropy in the string of twelve strategies each round) started to approach the random benchmark. However, individual entropies were diverging further below prediction during this same time period.

If groups eventually reach equilibrium, it may be that they do so despite their members. This would be consistent with patterns observed in the market entry game and the minority game (Bottazzi & Devetag, 2007; Duffy & Hopkins, 2005). These authors found that their participants had very predictable behavior, even after many iterations, and even though participants were well-compensated. And despite this individual-level predictability, group-level behavior in these experiments was indistinguishable from the randomness predicted by theory.

Group Effect

Participants collectively matched their behavior to that of their groups. This regularity is evident in a consistency in group members' rates of rotation. We propose that the consistency within groups is based on some unobserved function of other group members' inferred thinking steps. This reveals a new group-level effect of individual strategic reasoning—while partipants are in competition to earn points, they are also matching their rate of rotation to that of their peers.

At first glance, this claim about a group-level influence on reasoning may be mistaken for either of two more mundane claims. The first is the benign observation that participants and groups have individual differences. In the context of the Wheel Game, this explains behavior in the first rounds of play, but not the convergence of group members' rates of rotation over rounds. The observed convergence requires a more complex process unfolding from the initial conditions set by random variation.

The second related claim is that seeing others think ahead more will cause participants to think ahead more themselves. We support this finding from previous work, but our current claim is stronger. In the Wheel Game, seeing others think ahead causes partipants to think ahead by *a similar amount*.

We support both of these simple claims, but we also move beyond them. We provide evidence for a group-level effect of strategic thinking. Despite the variance between groups, a partipant's degree of strategic thinking is predicted by that of their groupmates. This degree differs by group, and this predictability increases over rounds as the members of a group implicitly coordinate.

Conclusion

Recognizing the limits of themselves and others to think entirely "around" a set of strategies, participants seem to have recruited the faculties behind the Beauty Pageant to a game with the structure of Matching Pennies. Investigating patterns in participants' rotation rates, we find evidence for a dynamic process of individual adaptation to an emergent influence on strategic behavior. Though we can only speculate on the mechanism, our work reveals the influence of an interpersonal mechanism on individual mentalizing behavior.

Acknowledgments

This work was supported by National Science Foundation REESE grant 0910218. The authors also wish to acknowledge Jenna Norden.

References

Bottazzi, G., & Devetag, G. (2007). Competition and coordination in experimental minority games. *Journal of Evolutionary Economics*, *17*(3), 241–275.

Camerer, C. (2003). Behavioral game theory: Experiments in strategic interaction. *New Jersey: Princeton University Press*.

Camerer, C., Ho, T., & Chong, J. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, *119*(3), 861–898.

Camerer, C., Ho, T., & Chong, J.-K. (2002). A cognitive hierarchy theory of one-shot games. *UCLA Department of Economics Working Paper*, 1–38. Available from `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=411061`

Camerer, C., & Hogarth, R. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, *19*(1), 7–42.

Costa-Gomes, M., Crawford, V., & Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, *69*(5), 1193–1235.

Duffy, J., & Hopkins, E. (2005). Learning, information, and sorting in market entry games: theory and evidence. *Games and Economic Behavior*, *51*(1), 31–62.

Goeree, J., & Holt, C. (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, *91*(5), 1402–1422.

Ho, T., Camerer, C., & Weigelt, K. (1998). Iterated dominance and iterated best response in experimental" p-beauty contests". *American Economic Review*, *88*(4), 947–969.

Ho, T., Camerer, C., & Weigelt, K. (2003). unpublished data. *in Colin Camerer, Behavioral Game Theory, Princeton University Press*, 18.

McCabe, K., Mukherji, A., & Runkle, D. (2000). An experimental study of information and mixed-strategy play in the three-person matching-pennies game. *Economic Theory*, *15*(2), 421–462.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, *85*(5), 1313–1326.

Rapoport, A., Stein, W., Parco, J., & Nicholas, T. (2003). Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior*, *43*(2), 239–265.

Shannon, C. E., & Weaver, W. (1949). A mathematical theory of communication. *The University of Illinois Press*.