Iterated Dominance and Iterated Best Response in Experimental "p-Beauty Contests"

By Teck-Hua Ho, Colin Camerer, and Keith Weigelt*

Picture a thin country 1000 miles long, running north and south, like Chile. Several natural attractions are located at the northern tip of the country. Suppose each of n resort developers plans to locate a resort somewhere on the country's coast (and all spots are equally attractive). After all the resort locations are chosen, an airport will be built to serve tourists, at the average of all the locations *including* the natural attractions. Suppose most tourists visit all the resorts equally often, except for lazy tourists who visit only the resort closest to the airport; so the developer who locates closest to the airport gets a fixed bonus of extra visitors. Where should the developer locate to be nearest to the airport?

The surprising game-theoretic answer is that *all* the developers should locate *exactly* where the natural attractions are. This answer requires at least one natural attraction at the northern tip, but does not depend on the fraction of lazy tourists or the number of developers (as long as there is more than one).[1]

[1] Douglas Gale (1995 Sec. 4) describes a related class of "dynamic coordination" games in which the returns to investing at time t depends on how much others invest *after* t. For example, a firm pioneering a new product standard benefits if more subsequent entrants use the same standard. In equilibrium, all firms invest immediately.

To see how this result comes about, denote developers' choices by mileage numbers on the coastline (from 0 to 1000) as $x_1, x_2, ... , x_n$. Locate all m of the natural attractions at 0. Then, the average location is $A = (x_1 + x_2 + \cdots + x_n)/(n + m) = n/(n + m)\cdot\bar{x}$. If we define the fraction $n/(n + m)$ as p (and note that $p < 1$ as long as $m \geq 1$), then the developer who is closest to A, or $p\cdot\bar{x}$, wins a fixed amount of extra business (from the lazy tourists).

This game was first discussed by Herve Moulin (1986 p. 72) and studied experimentally by Rosemarie Nagel (1995). It is solved by iterated application of dominance. The largest possible value of A is $1000\cdot p$ so any choice of x above $1000\cdot p$ is dominated by choosing $1000\cdot p$. If developers believe others obey dominance, and therefore choose $x_i < 1000\cdot p$, then the maximum A is $1000\cdot p^2$ so any choice larger than that is dominated. Iterated application of dominance yields the unique Nash equilibrium, which is for everyone to locate at zero. No matter where the average of the other developers' locations is, a developer wants to locate between that average and the natural attractions (which is where the airport will be built); this desire draws all the developers inexorably toward exactly where the attractions are.

We call these "p-beauty contest" games because they capture the importance of iterated reasoning John Maynard Keynes described in his famous analogy for stock market investment (as Nagel, 1995, pointed out). Keynes (1936 pp. 155–56) said

... professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole . . . . It is

not a case of choosing those which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees.

In Keynes's "newspaper competition" people want to choose *exactly* the same faces others choose. Professional investment is not quite like this. Think of the time at which investors choose to sell a rising stock as picking a number. When many investors choose to sell, the stock crashes; the time of the crash is around the average number (selling time) chosen. Then professional investment is a p-beauty contest (with $p < 1$) in which investors want to sell a few days ahead of the crash— picking a number equal to p times the average number—but not too far ahead.

Our paper reports experiments on p-beauty contest games. These games are ideal for studying an important question in game theory—how many iterations of dominance do players apply? The games are also useful for studying learning.

Our central contribution is application to p-beauty contests of a pair of structural models—a model of first-period choices, and a related model of learning—similar to those used recently by Debra Holt (1993), Dale Stahl and Paul Wilson (1994, 1995), and others. These models give a parsimonious way to empirically characterize the levels of iterated dominance and iterated best response. In the model of first-period choices, players are assumed to obey different levels of iterated dominance. We estimate from the data what the distribution of the different levels is most likely to be. The results show substantial proportions of players (at least 10 percent) obeying each of 0–3 levels. In the learning model, players are assumed to use different levels of iterated best response. A small fraction of players are adaptive, responding to past observations, but most exhibit some degree of sophistication, best responding to responses by adaptive learners.

I. Iterated Dominance and Iterated Best Response

A. *Iterated Dominance*

Iterated dominance is perhaps the most basic principle in game theory. Games in which iterated application of dominance determine a unique equilibrium are called "dominance solvable." In a dominance-solvable game, reaching the equilibrium requires some minimal number of steps of iterated dominance (which we call the rationality threshold).

The number of steps of iterated dominance people use is important for economics because many phenomena which appear irrational could be due to rational players expecting others to behave irrationally. That is, the "unravelling" toward unique equilibrium sometimes depends on many steps of iterated dominance. Potential examples include cooperation in finitely repeated prisoner's dilemma games, the winner's curse, escalating bids in the dollar auction, price bubbles in experimental markets (e.g., David Porter and Vernon L. Smith, 1995), and coordination games in which players must coordinate on the level of iterated rationality (see Ho and Weigelt, 1996).

There are good reasons to doubt that players behave as if they have more than a couple of steps of iterated rationality. Iterated reasoning is cognitively difficult. High levels of iterated rationality may not be easily justified by natural selection arguments (see Stahl, 1993). And many experimental studies indicate limited steps of iterated dominance.

Randolph Beard and Richard Beil (1994) found that about 50 percent of the subjects violated one step. Using a similar game, Andrew Schotter et al. (1994) found that half their subjects violated one step, and about 20 percent violated weak dominance in a matrix-form game. (The corresponding figures were only 8 percent and 2 percent in a tree-form game.) John Van Huyck et al. (1994) studied a five-strategy variant of a prisoner's dilemma, and found about three levels of iterated dominance. Camerer et al. (1996) studied the "electronic mail game" introduced by Ariel Rubinstein (1989) and observed about two levels of iterated dominance. In Richard D.

McKelvey and Thomas R. Palfrey's study (1992) of centipede games, most subjects revealed two or three levels of iterated dominance. At the risk of overgeneralizing across games which are too different, it seems that subjects rarely violate dominance, but usually stop after one–three levels of iterated dominance.

These studies did not try to carefully measure the fraction of subjects using various levels of iterated dominance. We use the method of Stahl and Wilson (1995) to do so. They define "level-0" players as those who choose strategies randomly and equally often, and "level-k" players as those who optimize against lower-level [level-$(k - 1)$] players. In three-strategy matrix games, they estimate that most players are level-1 or level-2.

p-beauty contest games are better than the games above for studying levels of iterated dominance for two reasons. First, p-beauty contests have many more strategies so they can detect violations of higher levels of iterated dominance.

Second, p-beauty contest games are constant sum. In experiments a test of dominance is usually a joint test of utility-maximization and the self-interest assumption (own-payoff-maximization). All the games described above are nonconstant sum games, in which violations of the self-interest assumption are more likely, and caistaken for dominance violations.[2] For example, in McKelvey and Palfrey's centipede games, 15–20 percent of the players who arrive at the final node violate dominance by "passing." Passing means taking 20 percent of a $32 pie (earning $6.40 while the other player gets $25.60) instead of 80 percent of a $16 pie (earning $12.80 while the other gets $3.20). Players who maximize utility but care enough about others' payoffs[3] will pass and appear to violate dominance. Since p-beauty contests

are constant sum, altruistic behavior will disappear as long as players care about their own payoffs more than those of others.

B. Learning and Iterated Best Response

The predominant view in modern game theory is that equilibria in all but the simplest games are reached by a learning or evolutionary process rather than by reasoning. Many processes have been studied (e.g., Paul Milgrom and John Roberts, 1991; Kenneth Binmore et al., 1995; Alvin Roth and Ido Erev, 1995; Camerer and Ho, 1998) but more careful empirical observations are needed to judge which rules describe learning best.

p-beauty contests are useful for studying learning empirically. Because convergence is not immediate, there is healthy variation in the data which can be used to study adaptive dynamics. And adaptive learners, who simply learn from past observations, choose different numbers than sophisticated learners who realize others are adapting; so the game can be used to estimate the proportions of adaptive and sophisticated types (Milgrom and Roberts, 1991).

Our experiments use several variants of the p-beauty contest game. First, we compare "finite-threshold" games (with $p > 1$) in which the equilibrium can be reached in a finite number of steps of iterated dominance, with "infinite-threshold" games with $p < 1$, in which the equilibrium cannot quite be reached in finitely many iterations of dominance. (The developer-location game in the introduction is an infinite-threshold game.)

Two other comparisons, between different group sizes n and values of p, are used to study whether the results are robust across parameter

---

[2] Another way to separate the two is to collect information-processing measures other than choices. For example, Camerer et al. (1993) show using measures of attention that players in a three-stage sequential bargaining game often do not look past the first stage, violating the computational underpinning of iterated dominance.

[3] For example, assume a linear additive "social utility function" in which player i's utility for the allocation (x_i,

x_j) is $x_i + \alpha x_j$ (traceable, at least, to Edgeworth). Then passing is utility-maximizing in centipede if and only if (iff) $\alpha > 0.29$. This point is even clearer from the large literature on ultimatum bargaining and public goods games. As many as half the players violate the conjunction of dominance and self-interest by rejecting low offers, giving money to others, contributing to a public good, or cooperating in a prisoner's dilemma (e.g., Camerer and Richard H. Thaler, 1995; John Ledyard, 1995; David Sally, 1995).

variations. Group size is particularly interesting because players in smaller groups exert more influence on the mean number; if they recognize this, they should choose lower numbers and converge more quickly, but in fact, small groups appear to converge more slowly.

C. Key Results

Our experiments and analysis yield several key results, which follow.

- First-period choices are widely distributed and far from equilibrium, but subsequent choices converge toward the equilibrium point (particularly in the finite-threshold game).
- First-period choices are consistent with a median of 2 steps of iterated dominance in the infinite-threshold game, and 1 step in the finite-threshold game. The estimated proportions are spread across levels 0–3 (at least 10 percent in each).
- Choices after the first period are consistent with 70 percent of the subjects best responding to a weighted sum of previous target numbers (weighting the previous target most strongly).
- The parameter estimates are sensitive to p, the group size, and whether subjects played a similar game before.

II. The p-Beauty Contest Game

In our experiments, a group of n subjects simultaneously choose a number from a closed interval $[L, H]$. The subject whose number is closest to p times the group mean wins $n \cdot \pi$. Thus, the expected payoff per subject is π even though n varies across groups. Denote the target number by $w = p \cdot \bar{x} = p \cdot (x_1 + x_2 + \cdots + x_n)/n$. Subjects' payoffs are determined as follows. Denote the set of winners I^* to be $\text{argmin}_i \{ | x_i - w | \}$ (the set of players whose choices are closest to w). Each winner $i \in I^*$ obtains a monetary prize of $n \cdot \pi / | I^* |$ (splitting $n \cdot \pi$ equally) and the remaining group members receive nothing. Variants of this game are denoted by $G([L, H], p, n)$.

Consider two variants of this game with n players: (1) finite threshold, $FT(p, n) =$ $G([100, 200], p, n)$, and (2) infinite threshold, $IT(p, n) = G([0, 100], p, n)$. While both games have an unique dominance-solvable equilibrium, $IT(p, n)$ requires an infinite level of iterated reasoning to solve the game, whereas $FT(p, n)$ requires only a finite level. Figures 1A–B illustrate this with $FT(1.3, n)$ and $IT(0.7, n)$. In these figures, the level of iterated rationality is indicated by $R(i)$. For example, $R(2)$ means that subjects are rational and know that others are rational. Figure 1A shows the threshold level needed to solve $FT(1.3, n)$ is 3. Subjects with zero levels of iterated rationality may choose numbers from $[100, 130)$ [i.e., $R(0)$[4]]. Rational players will choose a number from $[130, 200]$ because 130 dominates any number in $[100, 130)$ [i.e., $R(1)$]. (This illustration assumes a very large number of players, so that players can ignore their own effect on the mean and target number). Mutually rational players deduce it is in their interests to choose a number from $[169, 200]$ [i.e., $R(2)$]. To guarantee all subjects choose the unique equilibrium of 200 requires a threshold level of 3 [since a subject at $R(2)$ could choose less than 200]. When $p = 1.1$, the threshold level is 8.

In Figure 1B, the threshold level for the $IT(0.7, n)$ game is infinite. Rational players will only choose a number from $[0, 70]$ because any number in $(70, 100]$ is dominated by 70 [i.e., $R(1)$]. Applying the same reasoning, mutually rational $R(2)$ players will only choose a number from $[0, 49]$ (i.e., $0.7 \cdot 70$). With k-levels of iterated rationality, players will pick numbers from $[0, 0.7^{k+1} \cdot 100]$. All players will choose the unique equilibrium of 0 only if iterated rationality is infinite (i.e., $0.7^k \to 0$ as $k \to \infty$).

III. Experimental Design

To investigate the degree of iterated rationality we studied games with infinite (IT) and finite (FT) thresholds, and different values of

---

[4] Irrational players may also choose a number outside of $[100, 130)$ by chance; thus the number of players choosing between $[100, 130)$ is a lower bound on the number of $R(0)$ players.
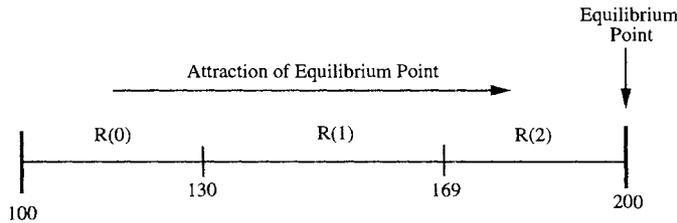
FIGURE 1A. A FINITE-THRESHOLD GAME, $FT(n) = ([100, 200], 1.3, n)$
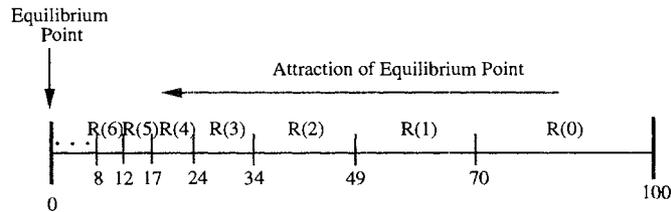
FIGURE 1B. AN INFINITE-THRESHOLD GAME, $IT(n) = ([0, 100], 0.7, n)$

p. The design also varied the group size n, to test whether smaller or larger groups behave differently. Table 1 summarizes the experimental design. Each subject played one IT game and one FT game (counterbalanced for order).

There were 55 experimental groups with a total of 277 subjects—27 groups of size 3, and 28 groups of size 7. Subjects were recruited from a business quantitative methods class at a major undergraduate university in Southeast Asia. They were assigned to experimental sessions randomly. Each participated in one session.

A typical session was conducted as follows. Subjects reported to a room with chairs placed around its perimeter, facing the wall, so subjects could not see the work of others. Subjects were randomly assigned seats, subject numbers, and given written instructions.[5] After all subjects were seated, an administrator read the instructions aloud, and subjects were given the

opportunity to ask questions.[6] During the experiment, subjects were not allowed to communicate with each other. Before round 1 began, all subjects were publicly informed of the relevant number range $[L, H]$, and the value of p. Then round 1 began. Subjects chose a number and wrote it on a slip of paper. An administrator then collected the responses of all subjects, calculated the average number, and publicly announced the average and the target number (p times the average). Payoffs were then privately announced to each subject

---

[5] The one-page instructions are available from the authors. Data and results of further analyses are also available.

[6] We did not adhere strictly to the standard protocol in experimental economics, in which subjects are only allowed to ask questions *privately*, and questions of general interest are repeated publicly by the experimenter. This protocol prevents subjects from communicating information ad lib which can have powerful effects and inhibit proper replication. (For example, if a subject asked: "Isn't the solution for everyone to choose 0?" publicly, that question could have a large effect which would be uncontrolled across treatments.) While allowing public questions was a mistake, there was only one such question (asking whether we were going to announce the target number) and it does not seem to have any perceptible influence (comparing our results with Nagel's, for example).

TABLE 1—THE EXPERIMENTAL DESIGN

Group size	
3	7
Finite → Infinite	
FT(1.3, 3) → IT(0.7, 3)	FT(1.3, 7) → IT(0.7, 7)
(7 groups)	(7 groups)
FT(1.1, 3) → IT(0.9, 3)	FT(1.1, 7) → IT(0.9, 7)
(7 groups)	(7 groups)
Infinite → Finite	
IT(0.7, 3) → FT(1.3, 3)	IT(0.7, 7) → FT(1.3, 7)
(7 groups)	(7 groups)
IT(0.9, 3) → FT(1.1, 3)	IT(0.9, 7) → FT(1.1, 7)
(6 groups)	(7 groups)

for that round [the winner(s) received a positive payoff; all others received \$0].[7] Then the next round began. All rounds were identical, and the game lasted for 10 rounds.[8] After the tenth round was completed, subjects participated in a second 10-round game in the same group, but with different parameter values (see Table 1). After this second 10 rounds, subjects summed their earnings over all 20 rounds, and were paid their earnings in cash. Experiments lasted approximately 40 minutes. The value of π was 0.5 Singapore dollars, so subjects

earned on average 10.00 Singapore dollars (about 7.00 U.S. dollars at that time).[9]

IV. Basic Results

This section summarizes basic results. Later sections report more refined structural estimates of the number of levels of iterated dominance and iterated best response subject use.

RESULT 1: *First-period choices are far from equilibrium, and centered near the interval midpoint. Choices converge toward the equilibrium point over time.*

Figures 2A–H show histograms of the frequencies of choices by subjects in each of the 8 IT conditions. Only 2.2 percent of the subjects chose the equilibrium in the first period. Most first-period distributions are sprinkled around the interval midpoint. Choices converge toward the equilibrium point.

There are fewer visible differences across FT conditions so the 8 FT conditions are collapsed into two histograms in Figures 3A–B, aggregating across 3- and 7-person-groups. The first-period distributions are also sprinkled around the interval midpoint, but choices converge more rapidly than in IT games across periods. About half the choices are exactly at the equilibrium of 200.

Table 2 summarizes the degrees of iterated dominance suggested by number choices. The table has four panels, one for each value of p, adding 3- and 7-person groups and experience

---

[7] Only the average and target were announced publicly. Payoffs to each subject were told to them privately. An alternative design is to withhold payoff information entirely. The potential problem with telling subjects their payoffs is that wealth effects alter incentives (which may, for example, contribute to the surprising outlying "spoiler" responses of 100). In Section IV we point out that there is no systematic evidence that these spoiler choices came from subjects who were satiated in money from winning repeatedly. Furthermore, one reason for privately informing subjects of their payoffs is that some learning models assume players know their own payoff history and only reinforce choices that they pick, and we wanted to generate data that could be used to test these models. If they were only told the target number, but not the payoff from the number they chose, these models would not apply and could not be tested.

[8] Each subject played the game with the same group members for all 10 rounds. Since the game was constant sum, there was no reason for subjects to develop reputation or tacitly collude to increase overall payoffs.

[9] The standard deviations of payoffs are 3.22 for $IT(p, 7)$, 2.48 for $FT(p, 7)$, 2.12 for $IT(p, 3)$, and 1.95 for $FT(p, 3)$. If subjects are equally skilled, the theoretical standard deviations are 3.89 ($n = 7$) and 2.24 ($n = 3$). Note that there is less variation in actual payoffs than predicted by the equal-skill benchmark. Most of the difference is due to the fact that subjects shared the prize in the event of a tie. Simulating the standard deviation of actual payoffs that would result if the whole prize were given to a randomly chosen subject in the event of a tie yields standard deviations that are extremely close to the equal-skill benchmark. Note that this finding casts some doubt on models in which players have persistent differences in skill, effort, or reasoning ability, etc., that create payoff differences. Or the standard deviation of payoffs may not be sensitive enough to detect individual differences.

FIGURE 2A. INEXPERIENCED SUBJECTS' CHOICES OVER ROUND IN $IT(0.7, 7)$

FIGURE 2B. EXPERIENCED SUBJECTS' CHOICES OVER ROUND IN $IT(0.7, 7)$

FIGURE 2C. INEXPERIENCED SUBJECTS' CHOICES OVER ROUND IN $IT(0.9, 7)$

FIGURE 2D. EXPERIENCED SUBJECTS' CHOICES OVER ROUND IN $IT(0.9, 7)$

FIGURE 2E. INEXPERIENCED SUBJECTS' CHOICES OVER ROUND IN $IT(0.7, 3)$

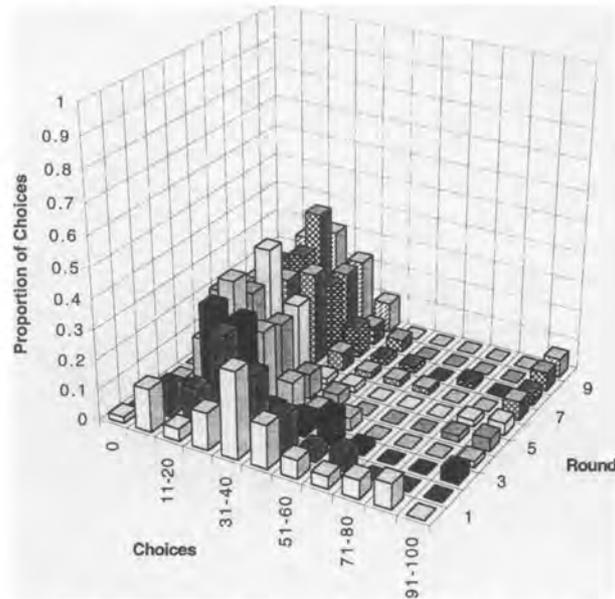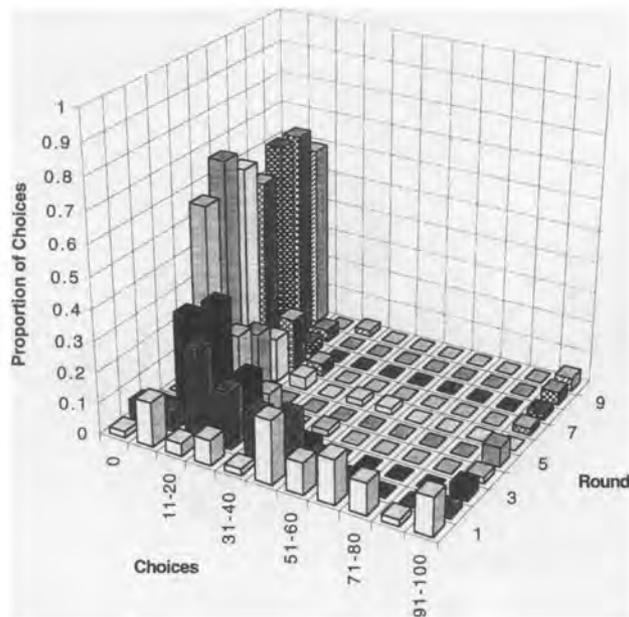FIGURE 2F. EXPERIENCED SUBJECTS' CHOICES OVER ROUND IN $IT(0.7, 3)$
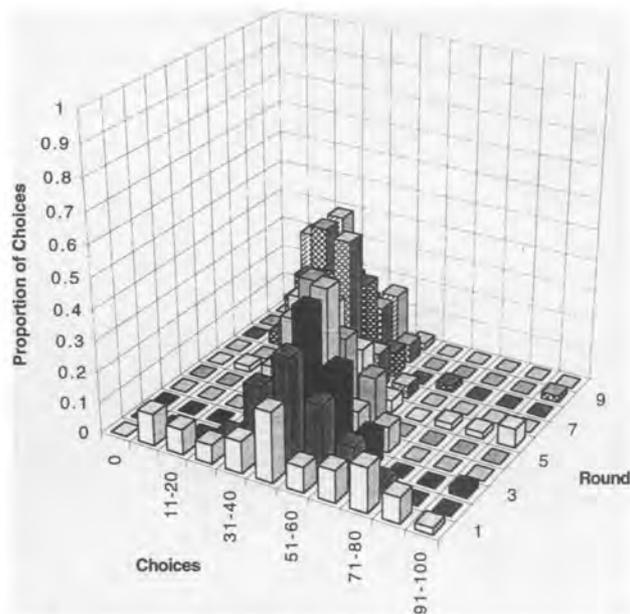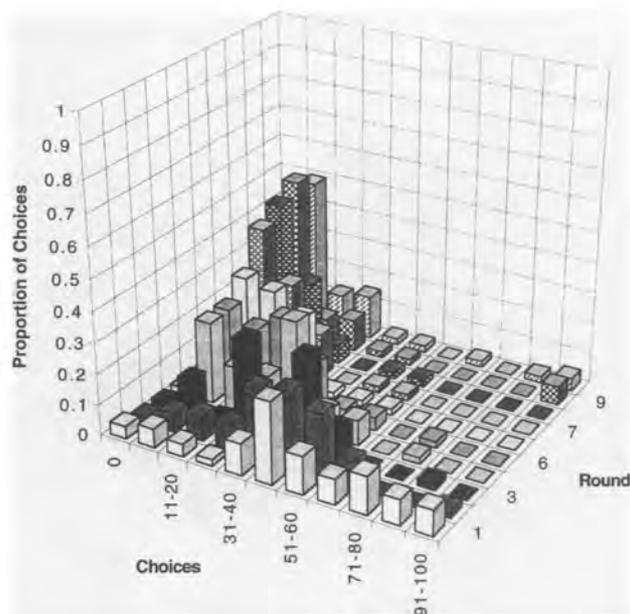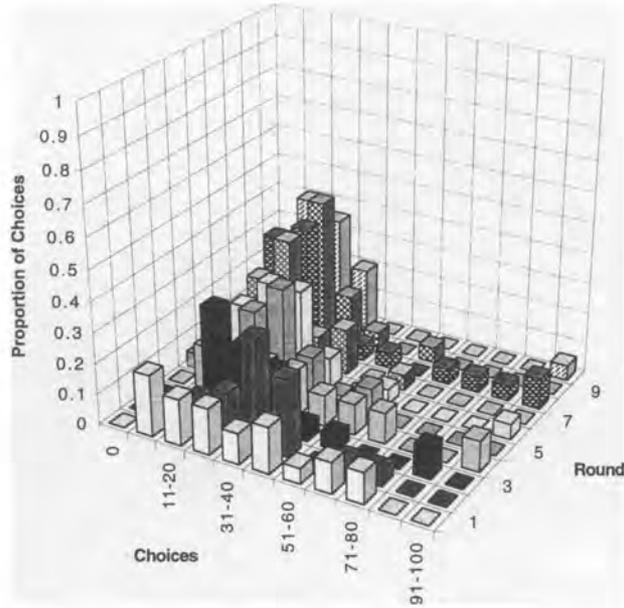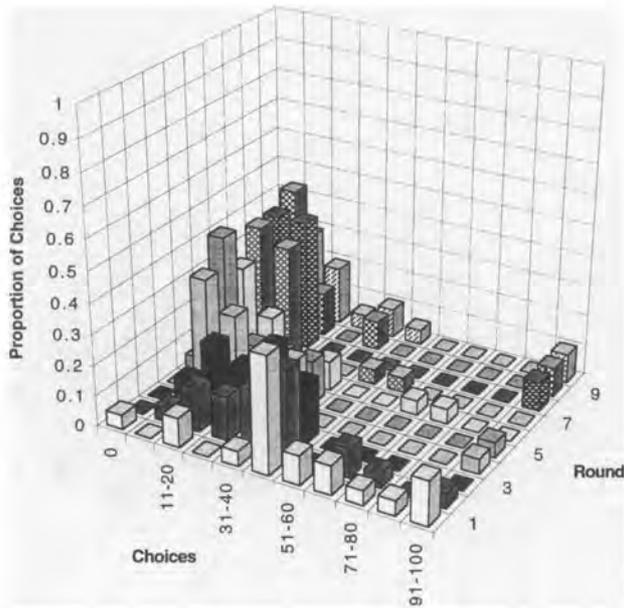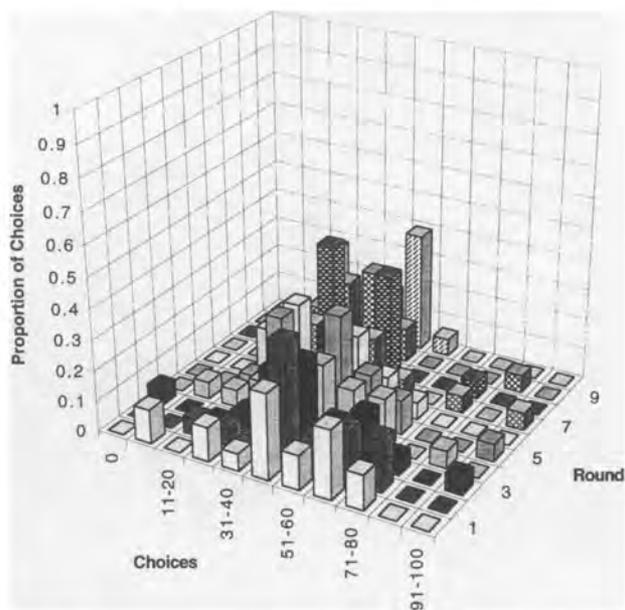
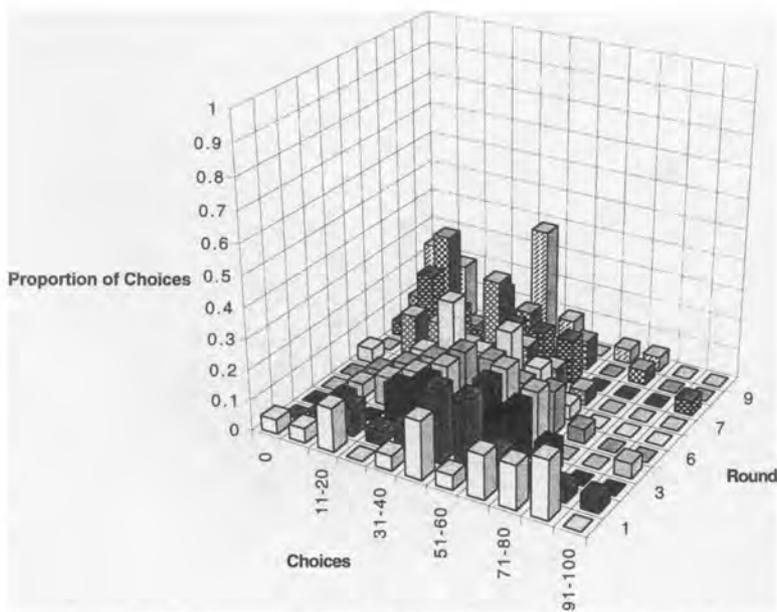FIGURE 2G. INEXPERIENCED SUBJECTS' CHOICES OVER ROUND IN $IT(0.9, 3)$

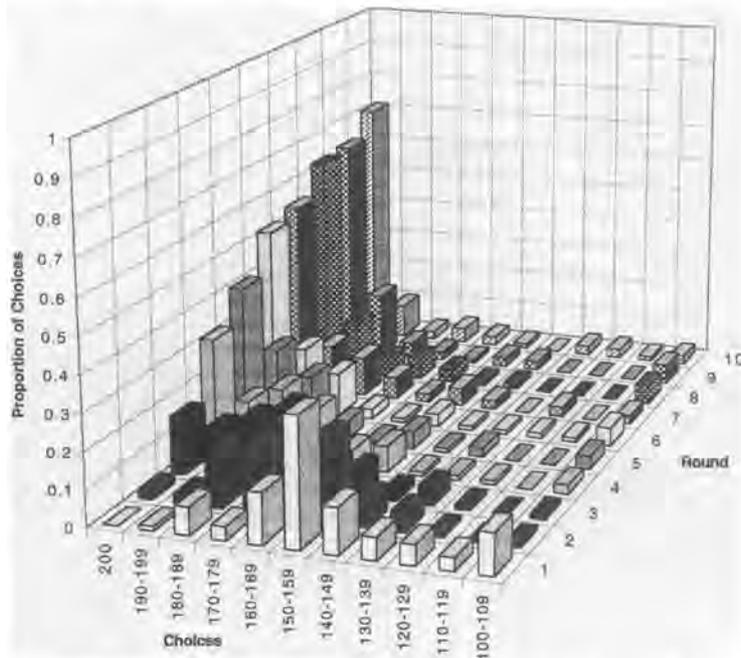FIGURE 2H. EXPERIENCED SUBJECTS' CHOICES OVER ROUND IN $IT(0.9, 3)$

FIGURE 3A. CHOICES OVER ROUND IN *FT* GAMES PLAYED BY 3-PERSON GROUPS

levels together for each p. Each row reports the number of choices violating each level of iterated dominance (in conjunction with lower levels). The last row of each panel reports the number of players who chose the unique equilibrium prediction.[10] For example, in condition $FT(1.3, n)$ in rounds 1–2, 44 responses out of 280 (15.7 percent) exhibited zero levels of iterated dominance because they fall in the interval [100, 130), 102 responses (36.4 percent) exhibited only one level of iterated dominance by choosing in [130, 169), and so forth.

Table 2 shows that substantial numbers of subjects exhibit each of the lowest levels of iterated dominance (and consequently, very few choose the equilibrium), particularly in earlier rounds. (Later rounds are consistent with higher levels of iterated dominance, but that is probably due to learning rather than

more sophisticated iterated reasoning per se.) Section V below gives more precise estimates.

RESULT 2: *On average, choices are closer to the equilibrium point for games with finite thresholds, and for games with p farther from 1.*

Comparing the Figure 2 and 3 histograms shows that after the first round, choices are further from equilibrium in the infinite-threshold games (Figure 2), compared to finite-threshold games (Figure 3).[11] The overall frequencies of equilibrium play are highly significantly different in *FT* and *IT* games (51.6 percent vs. 4.9 percent, $\chi^2 = 1493$, $p < 0.001$).

In addition, more choices are at equilibrium in games with p farther from zero ($p = 1.3$ vs. $p = 1.1$, $\chi^2 = 171.6$, $p < 0.001$; and $p = 0.7$ vs. $p = 0.9$, $\chi^2 = 5.9$, $p < 0.05$). Analyses of

---

[10] Levels of rationality higher than the threshold are indistinguishable, and pooled in "Equilibrium Play."

[11] Note that the tall bars in the back corner of Figures 2B and 2D represent frequent choices of 1–10, not zero.
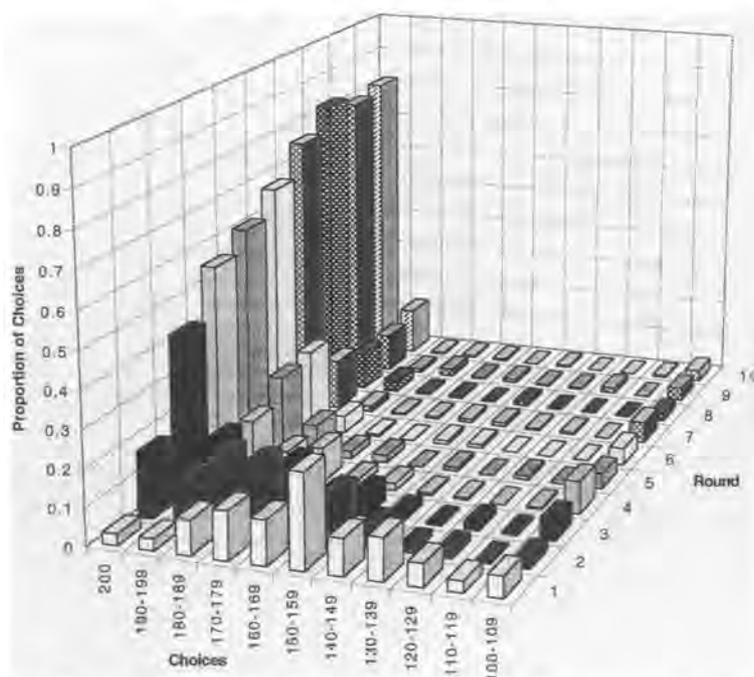
FIGURE 3B. CHOICES OVER ROUND IN *FT* GAMES PLAYED BY 7-PERSON GROUPS

variance (ANOVA) using each group's mean choice for the first or last five rounds also shows a strong effect of p, at significance levels from 0.000 to 0.08, for each of the two groups of rounds (1–5 or 6–10) and threshold levels. For example, means for $p = 0.7$ are lower than means for $p = 0.9$ for the first five rounds (31.57 vs. 44.66) and the last five rounds (17.76 vs. 27.83).

RESULT 3: *Choices are closer to equilibrium for large (7-person) groups than for small (3-person) groups.*

Figures 3A–B illustrate the typical effect of group size: Larger groups (Figure 3B) choose higher numbers at the start, and converge to equilibrium more quickly than the small groups in the same condition (Figure 3A). The quicker convergence in large groups is also evident in *IT* games (e.g., Figures 2A and 2E). Across rounds, the proportions of equilibrium play by subjects in 3- and 7-person groups are significantly different in finite-threshold games (39.6 percent vs. 56.6 percent, $\chi^2 = 67.3$, $p < 0.001$) and marginally significant in infinite-threshold games (3.7 percent vs. 5.4 percent, $\chi^2 = 3.4$, $p < 0.10$). Means of larger groups also start closer to the equilibrium for the *FT* games and, for both kinds of games, lie closer to the equilibrium in *every* round. ANOVAs on group means aggregated over rounds 1–5 or 6–10 show highly significant differences across group sizes in comparisons for all values of p and experience levels (p-values range from 0.003 to 0.014).

The group-size effect goes in a surprising direction because each member of a small group has a larger influence on the mean and should choose closer to equilibrium if he/she takes account of this. For example, if $p = 0.7$ and you think others will choose an average of 50, you should choose the solution to $C = 0.7 \cdot (C + (n - 1) \cdot 50)/n$, which is 30.4 if $n = 3$ and 33.3 if $n = 7$. But large groups

TABLE 2—FREQUENCIES OF LEVELS OF ITERATED DOMINANCE OVER ROUND IN FT AND IT GAMES
WITH VARYING p-VALUES

Games/Round	1–2	3–4	5–6	7–8	9–10	Total
FT(1.3, n)						
R(0)	44	27	14	14	11	110
R(1)	102	18	12	10	4	146
R(2)	101	70	49	22	7	249
Equilibrium Play	33	165	205	234	258	895
FT(1.1, n)						
R(0)	12	9	10	7	13	51
R(1)	9	2	4	2	3	20
R(2)	14	4	2	1	1	22
R(3)	27	7	5	4	2	45
R(4)	96	24	1	6	4	131
R(5)	65	59	13	7	11	155
R(6)–R(10)	42	103	118	76	72	411
Equilibrium Play	9	66	121	171	168	535
IT(0.7, n)						
R(0)	42	11	13	16	15	97
R(1)	65	21	5	7	3	101
R(2)	53	30	14	8	12	117
R(3)	35	53	37	21	21	167
R(4)	39	50	44	47	41	221
R(5)	13	43	35	36	32	159
R(6)–R(10)	25	71	108	102	91	397
>R(11)	2	1	12	18	25	58
Equilibrium Play	6	0	12	25	40	83
IT(0.9, n)						
R(0)	12	3	4	2	7	28
R(1)	7	2	1	0	1	11
R(2)	23	4	3	2	1	33
R(3)	17	12	1	0	2	32
R(4)	33	18	10	5	3	69
R(5)	14	21	12	6	3	56
R(6)–R(10)	117	142	100	80	60	499
>R(11)	47	69	136	162	175	589
Equilibrium Play	4	3	7	17	22	53

choose *lower* numbers. Perhaps, as a referee suggested, adjusting for n takes extra thought which limits the number of steps of iterated reasoning subjects do. This represents an interesting puzzle for future research.

RESULT 4: *Choices by experienced subjects are no different than choices by inexperienced subjects in the first round, but converge faster to equilibrium.*

We define experience to be cross-game experience, previous play with a *different p*-beauty contest (e.g., experience with an *IT* game for *FT*-game players, or vice versa). Obviously this kind of experience may differ

from repeated experience with the same game, but our paper does not explore this difference. Figures 2A–B illustrate the effects of experience in *IT* games. Experienced subjects choose similar numbers to inexperienced subjects in the first round, but converge much faster.

Two types of learning transfer by experienced subjects can be distinguished: "immediate transfer" (if choices are closer to equilibrium in the first round), and "structural transfer" if convergence is faster across the 10 rounds.

In general, there is little immediate transfer because experienced subjects' choices in the first round are not much different than the

choices of inexperienced subjects.[12] But there is some evidence of structural transfer because ANOVAs show that group means of experienced subjects are closer to equilibrium in the first five rounds ($F = 4.60$, $p < 0.04$) and about the same in the last five rounds ($F = 0.07$, $p < 0.80$) for FT games, and closer to equilibrium in IT games ($F = 2.62$, $p < 0.11$, and $F = 15.67$, $p < 0.001$). And overall, the proportion of equilibrium play is significantly higher for experienced subjects ($\chi^2 = 28.2$ (FT) and $\chi^2 = 125.6$ (IT), both $p < 0.001$).

Finally, a minor but eye-catching feature of the data evident in Figures 2–3 is the occasional choice of extreme numbers like 100. These ''spoiler'' choices occur 2.5 percent of the time and rarely occur more than once per subject. Spoilers tend to follow previous losses, so they are probably due to frustration or to misguided attempts to win by single-handedly raising the mean dramatically. Spoilers are even more rare in games where the target number depends on the median instead of the mean (Camerer et al., 1997; John Duffy and Nagel, 1997). The analyses below exclude spoilers (an alternative is to include a proportion of ''level − 1'' types as in Stahl, 1996a). Including them makes little substantive difference (see Ho et al., 1996).

V. Further Results: Levels of Iterated Dominance

The figures in Table 2 use the simplest method for approximating the level of iterated dominance revealed by choices: Count the number of choices in each of the intervals [0, $p^{k+1} \cdot 100$] (or the corresponding intervals when $p > 1$). For example, since 4.4 percent of the subjects chose numbers in (90, 100] in $IT(0.9, n)$ games in rounds 1–2, then we can conclude that *at least* 4.4 percent violated dominance. Since 2.6 percent of the subjects chose numbers in (81, 90], we can conclude

that at least 2.6 percent of the subjects violate the conjunction of dominance, and one level of iterated dominance. (Or put differently, we can be sure these subjects are not performing two or more levels of iterated dominance.)

These numbers are simply a lower bound on the rates at which various levels of rationality are violated. The bounds cannot be tightened without using some method for distinguishing how many of the 2.6 percent subjects in the interval (81, 90], for example, are violating dominance and how many are obeying dominance but violating one-step iterated dominance.

The method we use posits a simple structural model of how dominance-violating level-0 players choose, assume that level-1 players obey dominance but believe that others are level-0 players, etc. This method was first used by Stahl and Wilson (1994, 1995).[13]

Begin with the assumption that level-0 players choose numbers randomly from a truncated normal density with mean μ and variance σ^2.[14] Level-L players are assumed to believe that all other players (besides themselves) choose from the level $L − 1$ density $B_{L-1}(x)$. Believing this, they mentally simulate $n − 1$ draws from the level $L − 1$ distribution and compute the average of those draws. (For reasons we explain below, assume they allow these draws to be correlated with correlation ρ.) Then they choose p times the average (including their own choice), giving a density that satisfies

$$(1) \qquad B_L = \frac{p}{n} \cdot \left(B_L + \sum_{k=2}^{n} B_{L-1}^k \right),$$

where B_{L-1}^k is the k-th draw from random variable B_{L-1}. This gives a random variable for level-L players' choice

$$(2) \qquad B_L = \frac{p}{n - p} \cdot \sum_{k=2}^{n} B_{L-1}^k.$$

---

[12] The only exception is the $p = 0.7$ case where experienced subjects are *further* from the equilibrium point, exhibiting ''negative transfer.'' For $p = 0.7$, $t = −2.14$, $p < 0.01$ for $n = 7$, and $t = −2.15$, $p < 0.01$ for $n = 3$.

[13] While having only integer-level types is restrictive, Stahl (1996b) reported that allowing noninteger-level types does not improve fit of a rule learning model to Nagel's original data.

[14] We also tried a uniform distribution over all possible number choices, but the uniform almost always fit worse than the normal.

Notice that since the level-0 density is truncated at 100 (for the $p < 1$ case), the level-1 density is automatically truncated at $[p \cdot (n - 1)/(n - p) \cdot 100]$; so level-1 players never violate dominance. Similarly, since the level-2 density is truncated at $[p \cdot (n - 1)/(n - p)]^2 \cdot 100$, level-2 types never violate dominance and never violate one step of iterated dominance.[15]

It is easy to show that the mean and variance of B_L obey the following recursive relationships:

$$(3) \qquad E(B_L) = \frac{p \cdot (n - 1)}{n - p} \cdot E(B_{L-1}),$$

$$(4) \quad \mathrm{Var}(B_L) = \frac{p^2}{(n - p)^2} \cdot \left[(n - 1) + 2\rho \right.$$
$$\left. \cdot \frac{(n - 1) \cdot (n - 2)}{2} \right] \cdot \mathrm{Var}(B_{L-1}).$$

An important feature of this model is that if $p < 1$, as the level L rises, the variance in the distribution of choices B_L falls[16] [because the term $p^2/(n - p)^2 \cdot [(n - 1) + 2\rho \cdot (n - 1) \cdot (n - 2)/2]$ is less than one]. The variance falls because the players are assumed to take an average of $n - 1$ other players' choices, which will have less variance than an individual choice. This implies that the level-1 players' density will be rather narrow, the level-2 players' density will be narrower still, and so forth. Allowing higher-level players to per-

ceive a nonzero correlation ρ among (simulated) choices by lower-level players slows down the rate of reduction in $\mathrm{Var}(B_L)$ with L, and turns out to fit the data much better than the restriction $\rho = 0$.

The assumptions above give a density of first-period choices by each of the level types. The crucial problem is how to "assign" a level type to players who choose numbers x_i that different types might choose. Take the $p = 0.7$ case as a clarifying example. Suppose a player chooses 63. This choice could come from a level-0 player or from a level-1 player. We assign this choice to a level-0 type iff a level-0 type is more likely to have made that choice than a level-1 type [i.e., iff $B_0(63) > B_1(63)$]. Note that this method seems, at first blush, to be biased in favor of finding higher-level types: Since the higher-level types always choose closer and closer to the equilibrium, a person who chooses a low number randomly (a student who picks her age in years, for example) will seem to be misclassified as a high-level type. This is not true, however, because the *variance* of the high-type distributions shrinks (sometimes dramatically, if the group size is large). As a result, depending on the parameter values, only players who choose in a narrow range of low numbers will be classified as high types. [Note that a level-2 type would never pick 63, i.e., $B_2(63) = 0$, and similarly for higher-level types.]

Put more formally, assume that a fraction ω_L of the players are of level L, and ω_L^b is the fraction of choices assigned to L in each level-of-dominance interval or "bin" b. The total proportion of level-L types is $\omega_L = \Sigma_{b=0}^{L_m} N_b \cdot \omega_L^b / N$ (where L_m is the maximum level estimated [3 in our analyses] and N_b is the number of observations in bin b). Define the observations x in bin b by those x which satisfy $[p(n - 1)/(n - p)]^{b+1} \cdot 100 < x \le [p(n - 1)/(n - p)]^b \cdot 100$. (There are a total of $L_m + 1$ bins.) For observations in bin b, the density function is

$$(5) \qquad B(x) = \sum_{L=0}^{b} \omega_L^b \cdot B_L(x).$$

[15] One criticism of this method is that it assumes all players think they are "smarter" (or reason more deeply) than others. While this is logically impossible, it is consistent with a large body of psychological evidence showing widespread overconfidence about relative ability (see, e.g., Camerer and Dan Lovallo, 1996). An alternative approach includes some degree of "self-consciousness": Level-L types to believe that a fraction ω_L of others are level-L types like themselves, then perhaps impose (or test) the rational expectations assumption that the perceived ω_L is close to the econometrician's best estimate, given the data (as in McKelvey and Palfrey, 1992, 1995, 1996).

[16] In general, the same thing can be said for the case if $p > 1$ as long as ρ is small. However, if ρ is close to 1, then one can have $\mathrm{Var}(B_L) > \mathrm{Var}(B_{L-1})$.

Of course, $\sum_{L=0}^{L_m} \omega_L = 1$. Then the log-likelihood of observing a sample (x_i, $i = 1, \ldots, N$) is:

$$(6) \quad LL_1(\mu_1, \sigma_1, \rho, \omega_L^b; b = 0, \ldots, L_m;$$

$$L = 0, \ldots, L_m) = \sum_{i=1}^{N} \text{Log}(B(x_i)).$$

The objective is to maximize the log-likelihood LL_1 by choosing μ_1, σ_1, ρ, ω_L^b; $b = 0, \ldots, L_m$, $L = 0, \ldots, L_m$.

The left columns of Table 3 report maximum-likelihood parameter estimates for *IT* and *FT* games, using only first-round data.

In both games the estimates of ω_L show substantial proportions, at least 12 percent, in all level categories from 0–3. (Higher-level types are included in level 3). The median level is two for *IT* games and one for *FT* games. The *IT* games also have a larger fraction of high-level types than *FT* games. The estimated correlation ρ is 1.00 in both cases. This implies that higher-level subjects are choosing much more variable numbers than would be predicted if they were simply best responding to an average of independent choices by others. Their behavior is consistent with players choosing against a ''representative-agent player'' or composite, neglecting variation in the sample mean.[17]

The rightmost columns of Table 3 show parameter estimates using Nagel's (1995) data with $p = 1/2$ and $p = 2/3$ in *IT* games. Estimates

of normalized level percentages (from her Figure 2) are shown in parentheses. Our estimates using her data suggest more level 0's and fewer level 2's and 3's than in our data.

Our method for estimating the proportions of level types and Nagel's method give somewhat different results. Our estimates suggest more level 0's and fewer level 2's and 3's than her estimates. Her method posits an n-step reasoning process which begins from a reference point, 50, for the numbers reported in Table 3. Our method estimates this starting point, instead, giving a level-0 distribution mean of $\hat{\mu} = 35.53$ for $p = 1/2$ and $\hat{\mu} = 52.23$ for $p = 2/3$. She then tests the theory by counting the frequencies of choices in number intervals corresponding to various reasoning levels. Our structural method, in contrast, uses all the data and assigns each observation to some level of reasoning (based on relative likelihood), giving a more complete picture. For example, her method does not classify the 20 percent of subjects who choose greater than 50 in the first round in the $p = 2/3$ game (normalizing the percentages she reported spreads the 20 percent evenly over level categories). Our method mostly classifies these 20 percent as level 0's and consequently, we estimate a much higher level-0 proportion (our 28 percent vs. her 13 percent). Also, by estimating a lower level-0 mean than Nagel assumed (35.53 vs. 50) in the $p = 1/2$ case, we count more players as level-0 types.

VI. Further Results: Levels of Iterated Best Response

In this section we estimate a class of learning models to understand the dynamic process by which choices change over rounds. This class of learning models posits various levels of ''iterated best response'' and applies the same basic ideas in the last section of levels of iterated dominance underlying first-round choices, but applies it to learning over rounds.

In the model, level-0 learners simply choose a weighted sum of target numbers in previous rounds. Level-1 learners assume all others are level-0 learners and best respond to anticipated choices by level-0 learners. Level-2 learners best respond to level-1 learners, and so forth.

---

[17] The correlation ρ might be a game-theoretic incarnation of the ''representativeness'' heuristic in statistical judgment (see, for example, Daniel Kahneman and Amos Tversky, 1982). People using representativeness judge likelihoods of samples by how well they represent a population or process. Representativeness inadvertently neglects other statistical properties like variation—in this case, higher-level players neglect the fact that independent draws tighten the variance of the average they best respond to. A related phenomenon has been observed in experiments on ''weak link'' coordination games, in which a player's payoff depends upon his action and the minimum action chosen by others. The densities of first-round choices in these games is strikingly similar across groups of different sizes (even though the chance of getting a low minimum rises sharply as the group grows), as if players represent all other players as a single composite.

TABLE 3—MAXIMUM-LIKELIHOOD ESTIMATES AND LOG-LIKELIHOODS FOR LEVELS
OF ITERATED DOMINANCE (FIRST-ROUND DATA ONLY)

Parameter estimates	Out data (groups of 3 or 7)		Nagel's data (groups of 16–18)	
	$IT(p, n)$	$FT(p, n)$	$IT(0.5, n)$	$IT(2/3, n)$
ω_0	15.93	21.72	45.83 (23.94)	28.36 (13.11)
ω_1	20.74	31.46	37.50 (29.58)	34.33 (44.26)
ω_2	13.53	12.73	16.67 (40.84)	37.31 (39.34)
ω_3	49.50	34.08	0.00 (5.63)	0.00 (3.28)
μ	70.13	100.50	35.53 (50.00)	52.23 (50.00)
σ	28.28	26.89	22.70	14.72
ρ	1.00	1.00	0.24	1.00
$-LL$	1128.29	1057.28	168.48	243.95

These levels capture the distinction between adaptive learning (responding only to previous observations) and sophisticated learning (best responding to anticipated play by others) which is discussed, among others, by Milgrom and Roberts (1991). Level-0 learners are adaptive; higher levels are sophisticated.

To express the model formally, denote subjects' choices at round $t \in \{1, 2, \ldots, 10\}$ by $x_1(t), x_2(t), \ldots, x_n(t)$. The target number at round t, $w(t) = p \cdot \bar{x}(t) = p \cdot (x_1(t) + x_2(t) + \cdots + x_n(t))/n$. Suppose a subject of level L forms a guess $G_L^j(t)$ about what another subject j will choose. Given $G_L^j(t)$, the subject chooses a best response to maximize his or her expected payoff. That is, the subject will choose $B_L(t)$ such that

$$(7) \qquad B_L(t) = p \cdot \frac{B_L(t) + \sum_{j=2}^{n} G_L^j(t)}{n}.$$

Or

$$(8) \qquad B_L(t) = \frac{p}{n - p} \cdot \sum_{j=2}^{n} G_L^j(t).$$

The guess of level-L subjects at time t is assumed to be the best response of level $L - 1$'s (hence the term "iterated best response"), i.e.,

$$(9) \qquad G_L^j(t) = B_{L-1}(t).$$

The level-0 subject j is assumed[18] to choose randomly from a normal density with mean $\mu(t)$ equal to a weighted sum of the R previous target numbers (where R corresponds to level of recall), and variance $\sigma^2(t)$. That is,

$$(10) \qquad \mu(t) = \sum_{s=1}^{R} \beta_s \cdot w(t - s).$$

The parameters β_s capture the influence of past target numbers on the current choice. In addition, the correlation between subject choices in any level is ρ (for the same reasons given in the previous section). The standard deviation $\sigma(t)$ is allowed to grow or decline as follows:

$$(11) \qquad \sigma(t) = \sigma \cdot e^{\gamma \cdot t}.$$

Assume that a fraction α_L of the players are level-L best responders, and compose a

---

[18] In an earlier draft we assumed level-0 types chose a weighted sum of previous target numbers *and* their own previous choices. However, the coefficients on previous choices were rarely significant so those terms were dropped from specification (10). Our earlier analysis also allowed only level-1 types and referees wisely coaxed us to do this more general analysis.

mixture of the underlying densities to form an overall density of number choices, B. Then $B(x)$ is given by:

$$(12) \qquad B(x) = \sum_{L=0}^{L_m} \alpha_L \cdot B_L(x),$$

where L_m is the highest level allowed (restricted to three in our estimates).

When $L_m = 1$, level-0 learners choose from a normal density with mean given by (10) above and level-1 learners choose from a normal density with mean given by $p \cdot (n - 1)/(n - p)$. $[\Sigma_{s=1}^{R} \beta_s \cdot w(t - s)]$. This implies that when p is further from 1, and n is small, choices will be closer to equilibrium [because the fraction $p \cdot (n - 1)/(n - p)$ is less (greater) than one for IT (FT) games]. Our experiment was designed to test these predictions.

In addition, variants of three familiar special cases are nested in the general model. These are Cournot dynamics (Antoine Augustin Cournot, 1838), a variant of fictitious play (George Brown, 1951), and a hybrid case in which previous observations are given geometrically declining weight.

1. *Modified Fictitious Play.* —Fictitious play learning rules assume that the probability of another player's future choice is best predicted by the empirical frequency of that choice in previous plays. A plausible variant of this applied to the p-beauty contest game is that subjects are all level-1 learners who expect all others to choose an equally weighted average of the numbers they chose in the past. Modified fictitious play can thus be tested by restricting all types to be level 1 ($\alpha_1 = 1$), and $\beta_1 = \cdots = \beta_R = \beta$. (A further restriction is $\beta = 1/R$ but, as we shall see, that is strongly rejected.)

2. *Geometric Weighted Average.* —Fictitious play weights all previous observations equally. A more plausible model assigns geometrically decreasing weights to older observations, then averages them. The declining weight model will fit learning better if subjects realize that choices come from a nonstationary distribution (or others are learning too), and therefore give more recent observations more

weight. In this model, all types are level 1 ($\alpha_1 = 1$) and $\beta_s = \beta^s$.

3. *Cournot Dynamics.* —Cournot best-response dynamics assumes that players guess others will repeat their most previous choices—i.e., all types are level 1 ($\alpha_1 = 1$) and $\beta_1 = 1, \beta_2, \ldots, \beta_R = 0$.

The log-likelihood of observing a sample of N subjects over a total of 10 periods is given by:

$$(13) \qquad LL_2 = \sum_{i=1}^{N} \cdot \sum_{s=1}^{10} \text{Log}(B(x_i(s))).$$

Some subtle issues arise in implementing the estimation. The standard method in estimating models with R lags is to exclude the first R rounds of data. Since we estimate models with R up to 3, this means discarding 30 percent of the data. We fix initial conditions of the model by estimating a hypothetical "initial target number" $w(0)$. Level-0 learners are assumed to act as if they had observed the target number $w(0)$, before making their first-round choice. Continuing along the same lines, to estimate the model with $R = 3$ we estimate hypothetical target numbers $w(-1)$ and $w(-2)$. This method fixes the initial conditions, uses all the data, and uses the same data for different R values so they can be fairly compared.

There is also a heteroskedasticity problem in the data because the variance of choices generally falls over time. The likelihood function is a product of many normal densities. Since normal densities include a term which is the reciprocal of the standard deviation, densities become much larger when variances fall. As a result, later-period data will have a large influence on the coefficient estimates. To correct for this, we transformed the choices $x_i(t)$ at time t by subtracting the sample mean ($m(t)$) and dividing by the sample standard deviation ($s(t)$), giving transformed choices of $x_i'(t) = (x_i(t) - m(t))/s(t)$. If the sample size is large, the transformed choices at every period will have a standard normal density.[19]

---

[19] We also tried transforming the data by multiplying the choices in period t by $50/m(t - 1)$ in the IT games

The transformed data has a predicted mean of $(\mu(t) - m(t))/s(t)$ and standard deviation of $\sigma(t)/s(t)$. All results reported below are based on the transformed data.

Table 4 reports parameter estimates for recall lengths $R = 1$, 2 and 3.[20] Generally $R = 3$ fits best, as indicated by χ^2 statistics in the bottom rows comparing each model with the model with one less lag, though the improvement of fit over the one-period-lag $R = 1$ model is modest.

The estimates of learner-level proportions $\hat{\alpha}_L$ show that in both *IT* and *FT* games, a substantial portion of the players are level-1 best responders. Specifically, in *IT* games, about 70 percent of the players are level-1 learners and the remaining 30 percent are level 0's. In *FT* games, there are 11 percent, 68 percent, and 21 percent level 0's, 1's, and 2's, respectively. All models restricted to only one type are strongly rejected in favor of the many-type model.[21]

Estimates of the hypothetical initial target numbers $\hat{w}(0)$ are quite plausible, around 45 for *IT* games and 144 for *FT* games. A negative $\hat{\gamma}$ suggests that standard deviation declines over the periods, which is consistent with the observed data. The estimated initial standard deviations of level-0 choices, $\hat{\sigma}$, are reasonable too.

The estimates of $\hat{\beta}_1$ are above one for *IT* games and below one for *FT* games,[22] which

may seem odd since it implies that level-0 learners are picking numbers that are further from equilibrium than the previous target number. But keep in mind that level-1 learners choose a fraction $p \cdot (n - 1)/(n - p)$ of their guess about the average level-0 choice. When this fraction is multiplied by the typical $\hat{\beta}_1$, the product is usually around one or lower (for *IT* games), which captures the idea that level-1 players choose numbers around or lower than previous target numbers. If the estimates of $\hat{\beta}_1$ were much lower, that would force the level-1 choices to be "too small" to fit the data well.

We also estimated the learning-model parameters separately for each of the eight treatment combinations for both *IT* and *FT* games.[23] Parameter estimates are substantially different for different values of p as well as n and levels of experience, but not in an interesting way. The Cournot, fictitious play, and geometrically declining weight restrictions on β_s are all strongly rejected.[24] This is not surprising given the estimates in Table 4, because $\hat{\beta}_1$ is generally different from one and $\hat{\beta}_2$ as well as $\hat{\beta}_3$ are usually nonzero (rejecting Cournot). In addition, $\hat{\beta}$'s are different from each other and do not grow or decline geometrically (rejecting the other two theories).

Nagel (1996) reports informal tests of a "learning direction" theory. In learning direction theory, players are assumed to change their strategies in the direction of *ex post* best responses. Our working paper reports tests of two versions of this theory. One version is set-theoretic, and predicts the direction of change correctly 60 percent of the time. Another version, which is statistically comparable to the iterated best-response model, fits worse than the latter model (adjusted for degrees of freedom).

Stahl (1996a) proposes a kind of reinforcement learning model in which *decision rules* are reinforced rather than specific

---

(see Stahl, 1996a) and by reflecting choices around 100 and multiplying them by $50/(200 - m(t - 1))$ in the *FT* games. These transformations tended to "overscale" the data so that the variances of the transformed data in the later periods became much higher than those of the initial periods, when the later-period data are close to equilibrium.

[20] To check for robustness of transformation, we also estimated the parameters using the original data. The proportions of levels of iterated best response were essentially the same in every combination except for $R = 3$ in the *IT* games. In that combination, the proportions of levels 0 and 1 are 28.5 percent and 71.5 percent, respectively.

[21] The χ^2 statistics (twice the difference in log-likelihood) are 471 and 922 ($\alpha_0 = 1$ only), 1002 and 858 ($\alpha_1 = 1$), 1594 and 1012 ($\alpha_2 = 1$), and 2336 and 1124 ($\alpha_3 = 1$) for *IT* and *FT* games, respectively.

[22] The discrepancy in estimates $\hat{\beta}_s$ shows that parameters for *IT* and *FT* learning differ (the differing α_L estimates show this as well), so it can be rejected as a general theory of learning with invariant parameters.

[23] The eight analyses are not reported for the sake of brevity.

[24] For all data, using $R = 3$, the χ^2 statistics for Cournot, fictitious play, and geometric weights are 1252, 1780, and 1558 for *IT* games, and 1086, 1038, and 927 for *FT* games. Note that Cournot fits much worse than the $R = 1$ version because it forces $\beta_1 = 1$.

TABLE 4—MAXIMUM-LIKELIHOOD ESTIMATES AND LOG-LIKELIHOODS
FOR THE ITERATED BEST-RESPONSE LEARNING MODELS

Game parameter estimates	Infinite-threshold ($N = 2711$); Finite-threshold ($N = 2668$)		
	Recall period		
	$R = 1$	$R = 2$	$R = 3$
$IT(p, n)$			
α_0	0.2878	0.3132	0.2850
α_1	0.7122	0.6868	0.7150
α_2	0.0000	0.0000	0.0000
α_3	0.0000	0.0000	0.0000
β_1	0.962	1.464	1.414
β_2	—	−0.464	0.197
β_3	—	—	−0.573
$w(0)$	50.97	45.27	44.87
$w(-1)$	—	37.03	48.61
$w(-2)$	—	—	41.85
σ	38.66	30.122	41.08
γ	−0.118	−0.133	−0.125
ρ	0.000	0.000	0.000
LL	−2317.94	−2242.49	−2098.70
χ^2	—	150.90	287.58
$FT(p, n)$			
α_0	0.1185	0.1195	0.1135
α_1	0.6771	0.6801	0.6771
α_2	0.2044	0.2004	0.2094
α_3	0.0000	0.0000	0.0000
β_1	1.027	0.970	0.913
β_2	—	0.060	0.059
β_3	—	—	0.060
$w(0)$	149.08	148.13	143.657
$w(-1)$	—	154.64	222.159
$w(-2)$	—	—	224.626
σ	30.52	29.73	29.954
γ	−0.012	−0.008	−0.008
ρ	0.000	0.000	0.000
LL	−442.75	−437.80	−435.02
χ^2	—	9.90	5.62

strategies. In his rule-learning approach, subjects choose one of K decision rules, where decision rule k chooses a number equal to $p^k \cdot w(t - 1)$ after observing a previous target number $w(t - 1)$. Rules are reinforced by their expected payoff. Reinforcing rules, rather than specific numbers, *does* reinforce reasoning the "best" number of steps ahead. Fitting this model to Nagel's data, and using several other free parameters, Stahl finds that initial propensities toward k between 0–2 are about equal, and in about half the sessions propensities move toward $k = 2$ over four periods. In addition, he rejects a variety of alternative models (some nested, some not), including behavior reinforcement, forecasting of changes in the ratios of target numbers, and direction learning.

Two differences between the first-period estimation in Table 3, and the iterated best-response estimation in Table 4, are worth noting. First, our method estimates fewer high-level types (2 or 3) in the learning model. Since the learning model uses information from many rounds, there may be a natural tendency to classify players in mid-level types but we have not fully explored this possibility. (That is, if a player lurched from a level-0 type to a level-2 type across rounds, the data they generate might suggest existence of a level-1 type.) Another possibility is that subjects simply do more steps of iterated reasoning when thinking about the first round than when deciding how to react to experience. Second, the estimates of the perceived correlation between players' choices, $\hat{\rho}$, are around 1 for estimated iterated dominance and around 0 for estimated iterated best response. This is because there is more dispersion in the first-round data; to explain the broad range of data requires that the higher-level type distributions not shrink in variance too much, which requires estimating ρ is large (since distribution variance is increasing in ρ). Over rounds, however, actual choice variance shrinks so ρ is "allowed" to go toward zero and the model can still fit the data.

VII. Conclusion

Our results show that choices reveal a limited number of steps of iterated dominance [which could be taken as a sharp measure of the degree of bounded (mutual) rationality].

In the original work on these games, Nagel (1995) reports an average initial choice around 36, which corresponds to about two levels of iterated dominance. Duffy and Nagel (1997) basically replicated these patterns when the target number was the median, mean, or maximum chosen in a group. They find no substantial difference between mean and median games, and higher choices in the maximum game (see also Nagel, 1998).

Our results extend these earlier findings in several ways. We draw a novel distinction between games with finite and infinite rationality thresholds and show that finite-threshold games converge more quickly and reliably.[25] Our use of a Stahl-Wilson-type structural model of levels of reasoning, estimated from first-round choices, gives a sharper characterization of levels of iterated dominance. Using different group sizes reveals a puzzling effect— smaller groups learn slower. Playing different games sequentially shows some evidence of positive learning transfer.

By using 10 rounds (and collecting eight times as much data), we get a fuller picture of learning. The extra data enable reliable estimates of learning models. Tests of various learning models suggest two stylized facts. First, the data are consistent with presence of adaptive (level-0) learners who simply respond to experience, and sophisticated (level-1 and higher) learners who best respond to lower-level learners. Thus, any learning model which hopes to describe well should include both types. Second, familiar learning models which are special cases of our approach, including Cournot best-response dynamics (which looks back only one period) and fictitious play, are clearly rejected.

REFERENCES

Beard, Randolph and Beil, Richard. "Do People Rely on the Self-interested Maximization of Others—An Experimental Test." *Management Science*, February 1994, *40*(2), pp. 252–62.

Binmore, Kenneth; Gale, John and Samuelson, Larry. "Learning to Be Imperfect: The Ultimatum Game." *Games and Economic Behavior*, January 1995, *8*(1), pp. 56–90.

Brown, George. "Iterative Solution of Games by Fictitious Play," in Tjalling C. Koopmans, ed., *Activity analysis of produc-*

[25] Nagel conducted three sessions with $p = \frac{4}{3}$ and a number interval [0, 100]. First-period data look roughly like a reflection of $p = \frac{2}{3}$ data around the midpoint of 50, although many choose 33 and 100. However, with choices in [0, 100] and $p > 1$, no numbers are ruled out by any level of iterated dominance (in contrast to our design). Also, 0 and 100 are both equilibria in her design, though 0 is not trembling-hand perfect, but 200 is the unique equilibrium in ours.

tion and allocation. New York: Wiley,1951, pp. 374–76.

Camerer, Colin; Blecherman, Barry and Goldstein, David. "Iterated Dominance and Learning in 'Electronic Mail Game' Experiments." Working paper, California Institute of Technology, 1996.

Camerer, Colin and Ho, Teck-Hua. "Experience Weighted Attraction Learning in Normal-Form Games." *Econometrica*, 1998 (forthcoming).

Camerer, Colin; Ho, Teck-Hua; Rhee, Hongjai and Weigelt, Keith. "Incentives and Learning in *p*-Beauty Contest Games." Working paper, California Institute of Technology, 1997.

Camerer, Colin F.; Johnson, Eric J.; Rymon, Talia and Sen, Sankar. "Cognition and Framings in Sequential Bargaining for Gains and Losses," in Ken Binmore, Alan Kirman, and Pietro Tani, eds., *Frontiers of game theory*. Cambridge, MA: MIT Press, 1993, pp. 27–47.

Camerer, Colin F. and Lovallo, Dan. "Overconfidence and Excess Business Entry: An Experimental Approach." Working Paper No. 975, California Institute of Technology, 1996.

Camerer, Colin and Thaler, Richard H. "Anomalies: Ultimatums, Dictators and Manners." *Journal of Economic Perspectives*, Spring 1995, *9*(2), pp. 209–19.

Cournot, Antoine Augustin. "Recherches sur les Principes Mathématiques de la Théorié des Richesses." 1838. Translated into English by N. T. Bacon as *Researches into the mathematical principles of the theory of wealth*. London: Hafner, 1960.

Duffy, John and Nagel, Rosemarie. "On the Robustness of Behaviour in Experimental 'Beauty Contest' Games." *Economic Journal*, November 1997, *107*(445), pp. 1684–700.

Gale, Douglas. "Dynamic Coordination Games." *Economic Theory*, January 1995, *5*(1), pp. 1–18.

Ho, Teck-Hua; Camerer, Colin and Weigelt, Keith. "Iterated Dominance and Iterated Best-Response in Experimental *p*-Beauty Contests." Working paper, UCLA, 1996.

Ho, Teck-Hua and Weigelt, Keith. "Task Complexity, Equilibrium Selection, and Learning: An Experimental Study." *Management Science*, May 1996, *42*(5), pp. 659–79.

Holt, Debra. "An Empirical Model of Strategic Choice with an Application to Coordination Games." Unpublished manuscript, Queen's University, Canada, 1993.

Kahneman, Daniel and Tversky, Amos. "Judgment Under Uncertainty: Heuristics and Biases," in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press, 1982, pp. 3–20.

Keynes, John Maynard. *The general theory of interest, employment and money*. London: Macmillan, 1936.

Ledyard, John. "Public Goods: A Survey of Experimental Research," in John Kagel and Alvin E. Roth, eds., *Handbook of experimental economics*. Princeton, NJ: Princeton University Press, 1995, pp. 111–94.

McKelvey, Richard D. and Palfrey, Thomas R. "An Experimental Study of the Centipede Game." *Econometrica*, July 1992, *60*(4), pp. 803–36.

_____. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, July 1995, *10*(1), pp. 6–38.

_____. "Quantal Response Equilibria for Extensive Form Games." Working Paper No. 947, California Institute of Technology, 1996.

Milgrom, Paul and Roberts, John. "Adaptive and Sophisticated Learning in Normal Form Games." *Games and Economic Behavior*, February 1991, *3*(1), pp. 82–100.

Moulin, Herve. *Game theory for the social sciences*. New York: New York University Press, 1986.

Nagel, Rosemarie. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review*, December 1995, *85*(5), pp. 1313–26.

_____. "A Survey of Research on Experimental Beauty-Contest Games," in David Budescu, Ido Erev, and Robert Zwick, eds., *Games and human behavior: Essays in honor of Amnon Rapoport*, 1998 (forthcoming).

Porter, David and Smith Vernon L. "Futures Contracting and Dividend Uncertainty in Experimental Asset Markets." *Journal of Business*, October 1995, *68*(4), pp. 509–42.

Roth, Alvin and Erev, Ido. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior*, Special Issue: Nobel Symposium, January 1995, *8*(1), pp. 164–212.

Rubinstein, Ariel. "The Electronic Mail Game: Strategic Behavior Under 'Almost Common Knowledge'." *American Economic Review*, June 1989, *79*(3), pp. 385–91.

Sally, David. "Conversation and Cooperation in Social Dilemmas: Experimental Evidence from 1958 to 1992." *Rationality and Society*, January 1995, *7*(1), pp. 58–92.

Schotter, Andrew; Weigelt, Keith and Wilson, Charles. "A Laboratory Investigation of Multiperson Rationality and Presentation Effects." *Games and Economic Behavior*, May 1994, *6*(3), pp. 445–68.

Stahl, Dale. "Evolution of $Smart_n$ Players." *Games and Economic Behavior*, October 1993, *5*(4), pp. 604–17.

_____. "Boundedly Rational Rule Learning in a Guessing Game." *Games and Economic Behavior*, October 1996a, *16*(2), pp. 303–30.

_____. "Is Step-j Thinking an Arbitrary Modelling Restriction or a Fact of Human Nature?" Unpublished manuscript, University of Texas, 1996b.

Stahl, Dale and Wilson, Paul. "Experimental Evidence on Players' Models of Other Players." *Journal of Economic Behavior and Organization*, December 1994, *25*(3), pp. 309–27.

_____. "On Players Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior*, July 1995, *10*(1), pp. 218–54.

Van Huyck, John; Wildenthal, John and Battalio, Raymond. "Tacit Cooperation, Strategic Uncertainty, and Coordination Failure: Evidence from Repeated Dominance Solvable Games." Working paper, University of Texas A&M, 1994.