Task Complexity, Equilibrium Selection, and Learning: An Experimental Study

Teck-Hua Ho • Keith Weigelt
Anderson Graduate School of Management, University of California, Los Angeles, California 90095
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104

We consider several coordination games with multiple equilibria each of which is a different division of a fixed pie. Laboratory experiments are conducted to address whether "task complexity" affects the selection of equilibrium by subjects. Three measures of task complexity—cardinality of choice space, level of iterative knowledge of rationality, and level of iterative knowledge of strategy—are manipulated and tested. Results suggest the three measures can predict choice behavior. Since strategically equivalent games can have different task complexity measures, our results imply that subjects are sensitive to game form presentation. We also fit data using three adaptive learning models: 1) Cournot, 2) Fictitious Play, and 3) Payoff Reinforcement, in increasing order of required cognitive effort. The Fictitious Play model, which tracks only cumulative frequencies of opponents' past behaviors fits the data best.
(*Task Complexity; Equilibrium Selection; Learning: An Experimental Study*)

1. Introduction

A challenging and persistent issue in game theory is multiplicity of equilibria: in games with multiple pareto efficient equilibria, there are no systematic criteria for equilibrium selection. Games with multiple equilibria are common; examples include signaling, coordination, and infinitely repeated games. Game theorists have tried to predict choice behavior in such games by modeling the degree of rationality and sophistication of players (see Myerson 1992, Fudenberg and Tirole 1992, Binmore 1992 for reviews). Starting with Selten (1965), the "refinement" literature addresses the multiplicity problem by imposing greater reasoning sophistication on players. This sophistication goes beyond the standard Nash rationality assumption that a player will choose his optimal response given others are choosing theirs. A sophisticated player will not make or believe incredible threats or promises, and will not draw unbelievable inferences from information she gathers (Kreps 1990a). The refinement approach has successfully prescribed a unique equilibrium in some signaling games (Cho and Kreps 1987, Banks and Sobel 1987), but is less successful in both infinitely repeated and coordination games.

More recently, another stream of literature imposes a lesser degree of rationality on players. Some believe the multiplicity of equilibria problem may be better attacked via this "bounded rationality" approach (Lucas 1986, Kreps 1990b). Players are bounded in several ways—they are myopic; they update beliefs using specific models (e.g., Cournot, fictitious play); they use different choice rules (e.g., SEU, imitation); or their behavior resembles programs in computing machines (e.g., finite automata, Turing machines). This approach approximates human choice behavior better, and has shown promise. Rubinstein (1986) shows the equilibria set in infinitely repeated prisoner's dilemma games is smaller if players act as finite automata. Marimon, McGrattan and Sargent (1990) show that even if players are modeled as simple computing machines, they will converge to a Nash equilibrium behavior in a complex macroeconomic setting. An extreme form of bounded rationality is captured in evolutionary stable strategies (ESS). Here, players are modeled as pre-programmed "types" where a player's type defines her strategy

0025-1909/96/4205/0659$01.25

(Maynard Smith 1982). For example, in a prisoner's dilemma game, a cooperative type will always play cooperatively no matter what strategies others follow.

The above two approaches vary the degree of rationality to solve the multiplicity of equilibria problem. A possible complementary approach is to model an equilibrium's task complexity. The idea is straightforward: If players are boundedly rational they may select equilibrium based on task complexity. An equilibrium's task complexity measures the degree of cognitive effort along the equilibrium path required to reach that equilibrium. Indeed, task complexity and bounded rationality are flip sides of a coin. If players are fully rational, task complexity is irrelevant because they can solve games of any complexity. Conversely, games with trivial complexity can be solved by even severely boundedly rational players. Thus, to address the multiplicity of equilibria problem, one must recognize this interdependency between bounded rationality and task complexity.

The degree of cognitive effort required to reach an equilibrium may differ among equilibria. It depends on the number of information sets, the number of alternatives at each information set, and the amount of necessary knowledge about others' behavior. If players are subject to information processing constraints (i.e., boundedly rational), they may adopt cognitive effort conserving rules, (i.e., they may choose equilibrium paths with low task complexity). Three task complexity metrics are more formally defined in §2. We test the effects of task complexity on equilibrium selection using laboratory experiments.[1]

Task complexity metrics are sensitive to presentation changes in strategically equivalent games in an interesting way. They are not invariant to transformations on extensive-form games that preserve the reduced normal-form (Thompson 1952, Dalkey 1953).[2] Equilibria concepts defined for reduced normal-form games (e.g., proper, persistence, stable sets) inherently assume players are insensitive to game form or simply restrict the domain of the theory's application to reduced normal form games. Conversely, equilibrium concepts for extensive-form games (e.g., sub-game perfect, sequential, Harsanyi-Selten 1 point solution) do allow sensitivity to game form.[3] Therefore, theoretically it is unclear whether a strategically equivalent game presented in different game forms should have different solution sets. Our task complexity metrics will provide evidence for this game form dependency debate. For if players are game form dependent they are also sensitive to task complexity metrics (because metrics are game form dependent).

We also investigate the interdependency between task complexity and bounded rationality by examining whether subjects' learning behavior is contingent on task complexity. We test three classical adaptive learning models to see whether they fit subjects' choice behavior across rounds. The first resembles a Cournot learning model; the second, a modified Fictitious Play model; and the third, a Payoff Reinforcement model.

This paper is organized as follows: §2 defines the indices of task complexity. Experimental design is explained in §3. Hypotheses are tested and results reported in §4. Implications and future research are discussed in §5.

2. Task Complexity Metrics

We consider extensive-form coordination games with 2 or 3 pareto efficient equilibria. The equilibria are alternative divisions of a fixed "pie." Equilibria differ along three task complexity measures. The three measures are:

Cardinality of choice space (c_x) of equilibrium x is formally defined as:

$$c_x = \max_y\{\text{no. of outcomes at stage } y \text{ along } x\text{'s equilibrium path}\}$$

where a stage is equivalent to an information set facing the player along the path leading to equilibrium x.[4]

[1] The focus on task complexity has another merit. It allows greater experimental control since it is easier to manipulate task complexity than the players' reasoning sophistication.

[2] A reduced normal-form game has multiple extensive form equivalents. Invariant transformations allow one to map an extensive-form game to its reduced normal form or to its other extensive-form equivalents (see Figure 8).

[3] For example, in Figure 8, game F1 has 2 sub-game perfect equilibria (2, 2 and 3, 3). The strategically equivalent game F2 has only one sub-game perfect equilibrium (3, 3).

[4] A reviewer aptly pointed out that an alternative measure is the average number of outcomes along an equilibrium path. Using this al-

Our metric is equivalent to the complexity measures used by Johnson and Payne (1985) in nonstrategic choice situations. Based on a series of Monte-Carlo simulations, they show that elementary cognitive operations increase linearly with the number of choice outcomes. This implies that elementary cognitive operations should increase with c_x. Consequently, if subjects adopt effort conserving rules that favor fewer elementary operations, their choice behaviors will be more consistent with choosing an equilibrium with a smaller c_x.

For example, in Figure 1, the (8, 4) equilibrium (bottom) has a c index of 4 because there are 4 possible outcomes in both Stages I and II. Whereas, the (7, 5) equilibrium (top) has a c index of 9 because there are 4 possible outcomes in Stage I, and 9 in Stage II.

The **level of iterative knowledge of rationality** (r_x) of equilibrium x is equal to the number of mutual rationality levels needed to reach equilibrium x. A level of mutual rationality is defined as each player knowing the other is an expected-utility maximizer; he knows the other's action will maximize her own expected utility (Brandenburger 1992). r_x represents the required level of conjectures or presumptions about whether others are rational or not in order to reach equilibrium x. In a normal-form game, this measure corresponds to the rounds of iterative dominance; in an extensive-form game this is equivalent to the number of backward induction steps.

Figure 2 illustrates this. Starting from the far right (i.e., backwardly inducting), a rational I will take at this stage because 16 is greater than 10. Moving to the left, if II knows that I is rational, a rational II compares a payoff of 4 (pass) and 8 (take), thus he takes. This requires a level of iterative knowledge of rationality of 1. Again moving to the left, if I knows 1) II is rational and 2) II knows that I is rational, then a rational I will take because 4 is greater than 2. This requires a level of iterative knowledge of rationality of 2.

ternative measure in our design would not affect results, since the more complex games (as measured by our maximum rule) are also more complex using the suggested average rule. We use the maximum rule because we believe subjects decompose a game into stages before solving it and we are curious about their cognitive limits. We leave to future experiments a test of which rule predicts behavior better.

Although Johnson and Payne (1985) only looked at non-strategic games, we can map our metric into their information processing model. As the level of iterative knowledge of rationality increases, the number of elementary cognitive operations necessarily increase. That is, as players proceed to a higher level of recursive reasoning, they must read, move, and compare the results of all conjectures at lower levels (i.e., total elementary operations increase with r_x).

The **level of iterative knowledge of strategies** (s_x) of equilibrium x is equal to the number of levels of mutual knowledge of strategies needed to reach equilibrium x. There is an important, but subtle, difference between the mutual knowledge of rationality and mutual knowledge of strategies. Instead of knowledge about the player's rationality, this knowledge pertains to a player's action choice; this is especially critical in the presence of multiple equilibria. It is not sufficient for players to know others are rational, they must know the strategy others will play.

Figure 3 illustrates this: For player I to take (7, 5), she must know that player II is also choosing equilibrium (7, 5). So the s index is 1. The s index for the other two equilibria (3, 9) and (8, 4) is 2. Here, both players must believe that the other will go to Stage II, and within Stage II each player will choose the appropriate strategy. Knowledge of strategies supersedes knowledge of rationality. If a player knows what strategy others will choose, she does not have to care whether others are rational.

The iterative knowledge of strategies of an equilibrium is similar to the number of stages along the path leading to the equilibrium. An equilibrium that has a higher mutual knowledge of strategies also has a higher number of stages (or information sets) along its equilibrium path. Similar to iterative knowledge of rationality, each level of knowledge of strategies requires players to perform additional elementary cognitive operations. Hence, equilibria with high s_x have a higher number of elementary operations.

We characterize each equilibrium by four indices: cardinality of choice space (c), level of iterative knowledge of rationality (r), level of iterative knowledge of strategies (s), and the level of payoff disparity (p). We control for payoff disparity because previous experiments using fixed-pie games suggest subjects are sensitive to it

Figure 1 Cardinality of Choice Space

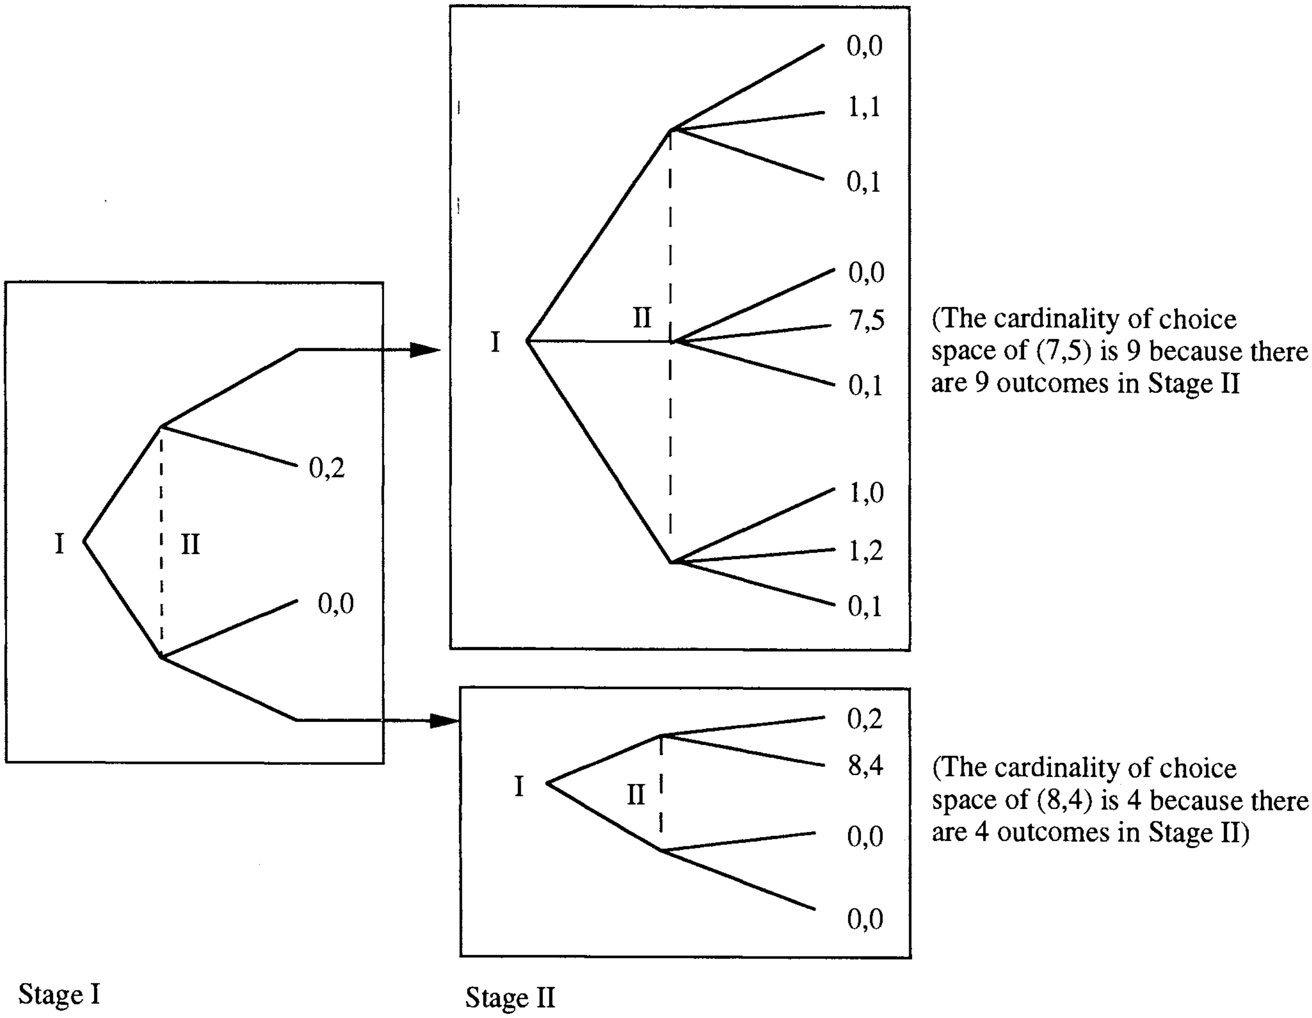

Figure 2 Level of Iterative Knowledge of Rationality

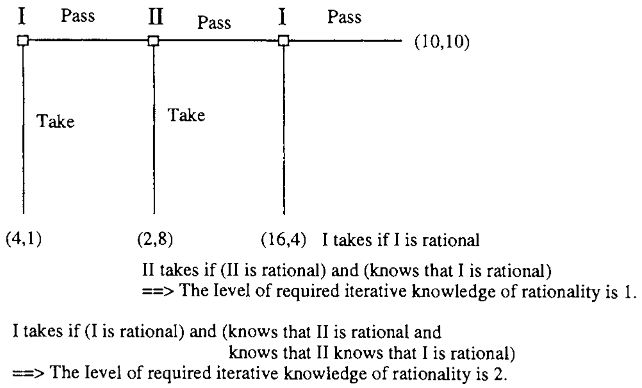

Figure 3 Level of Iterative Knowledge of Strategies

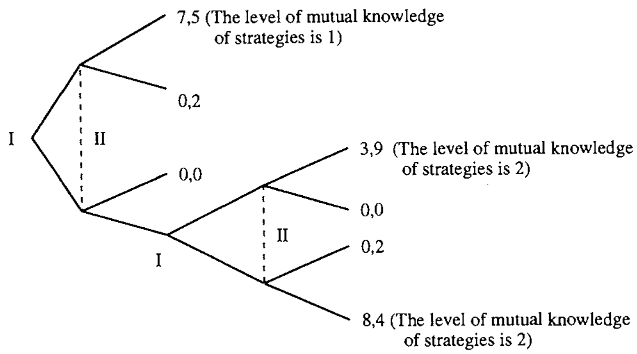

Figure 4 Experimental Design for Determining the Effect of Cardinality of Choice Space

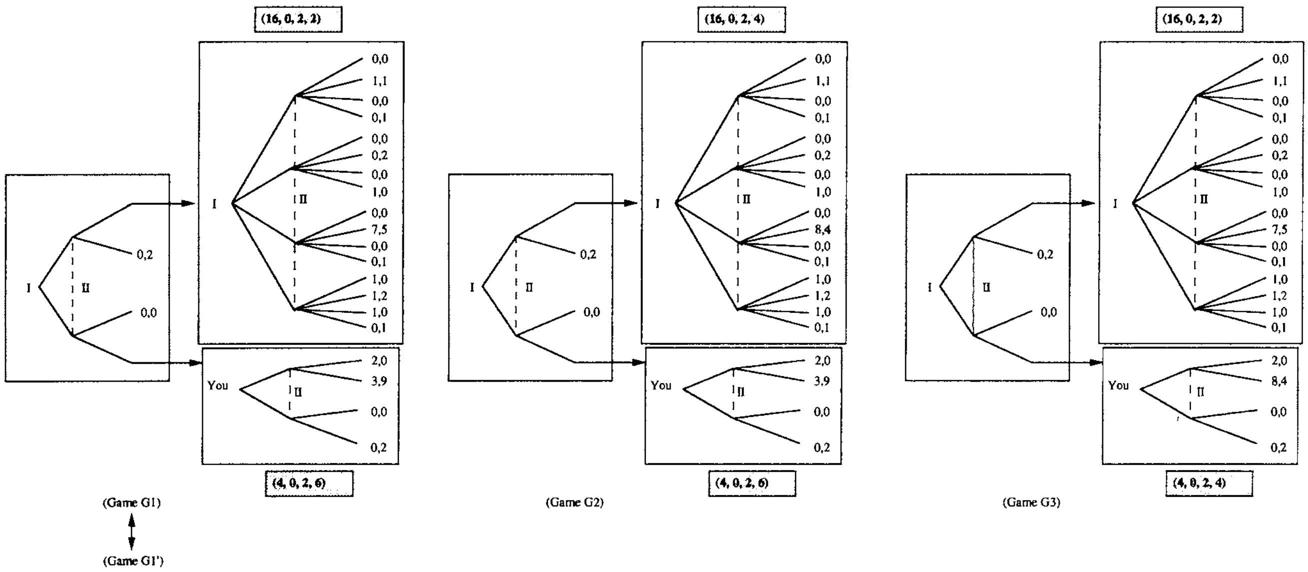

(see Roth, forthcoming).[5] This finding suggests payoff disparity as a possible equilibrium selection criterion, and offers an alternative hypothesis for observed behavior. We define the payoff disparity of equilibrium x (p_x) as the difference between the payoffs of the two players. For example, the p of equilibrium (7, 5) is 2. Hence, each equilibrium x is represented as a 4-dimensional vector (c_x, r_x, s_x, p_x). c_x, r_x, and s_x capture task complexity.

In summary, our task complexity metrics provide a proxy for the degree of cognitive effort required to play a game. Starting with the simplest decision problem, games against nature, only cardinality of choice enters the problem's task complexity, since it is a single-person choice problem. In games solved by iterative dominance or backward induction (i.e., games of perfect and complete information), cardinality of choice and iterative knowledge of rationality are sufficient (Figure 2), since there is only one node in each information set. Iterative knowledge of strategies is needed in more general classes of games, including; normal form pure strategy games, and extensive-form games where there is imperfect information (Figure 3) since there exists multinode information sets. Equilibria that have higher c, r, or s can only be reached if subject adopt rules which consist of more elementary cognitive operations. Consequently, if subjects adopt effort conserving rules, they will reach those equilibria with smaller complexity metrics.

[5] For example, in ultimatum games, subjects usually reject offers where they receive less than 20% of the fixed payoff. This behavior is irrational if players ignore payoff disparity.

3. Experimental Design and Procedure

3.1. Experimental Procedure

Subjects were recruited from both Wharton undergraduate and graduate classes. Students were randomly selected to meet at the behavioral laboratory at a specified time. The behavioral laboratory consisted of cubicles that made it difficult for subjects to see others during the experiment. Subjects were randomly assigned to a cubicle, and instructions were handed out (see Appendix) and read publicly. The instructions basically said the following: All subjects were to be given the same decision problem, and each subject would make a series of 10 decisions with a randomly selected "pair member." Subjects did not know the identity of their pair member. While the decision problem remained constant over the 10 rounds,

Figure 5 **Experimental Design for Determining the Effect of Iterative Knowledge of Rationality**

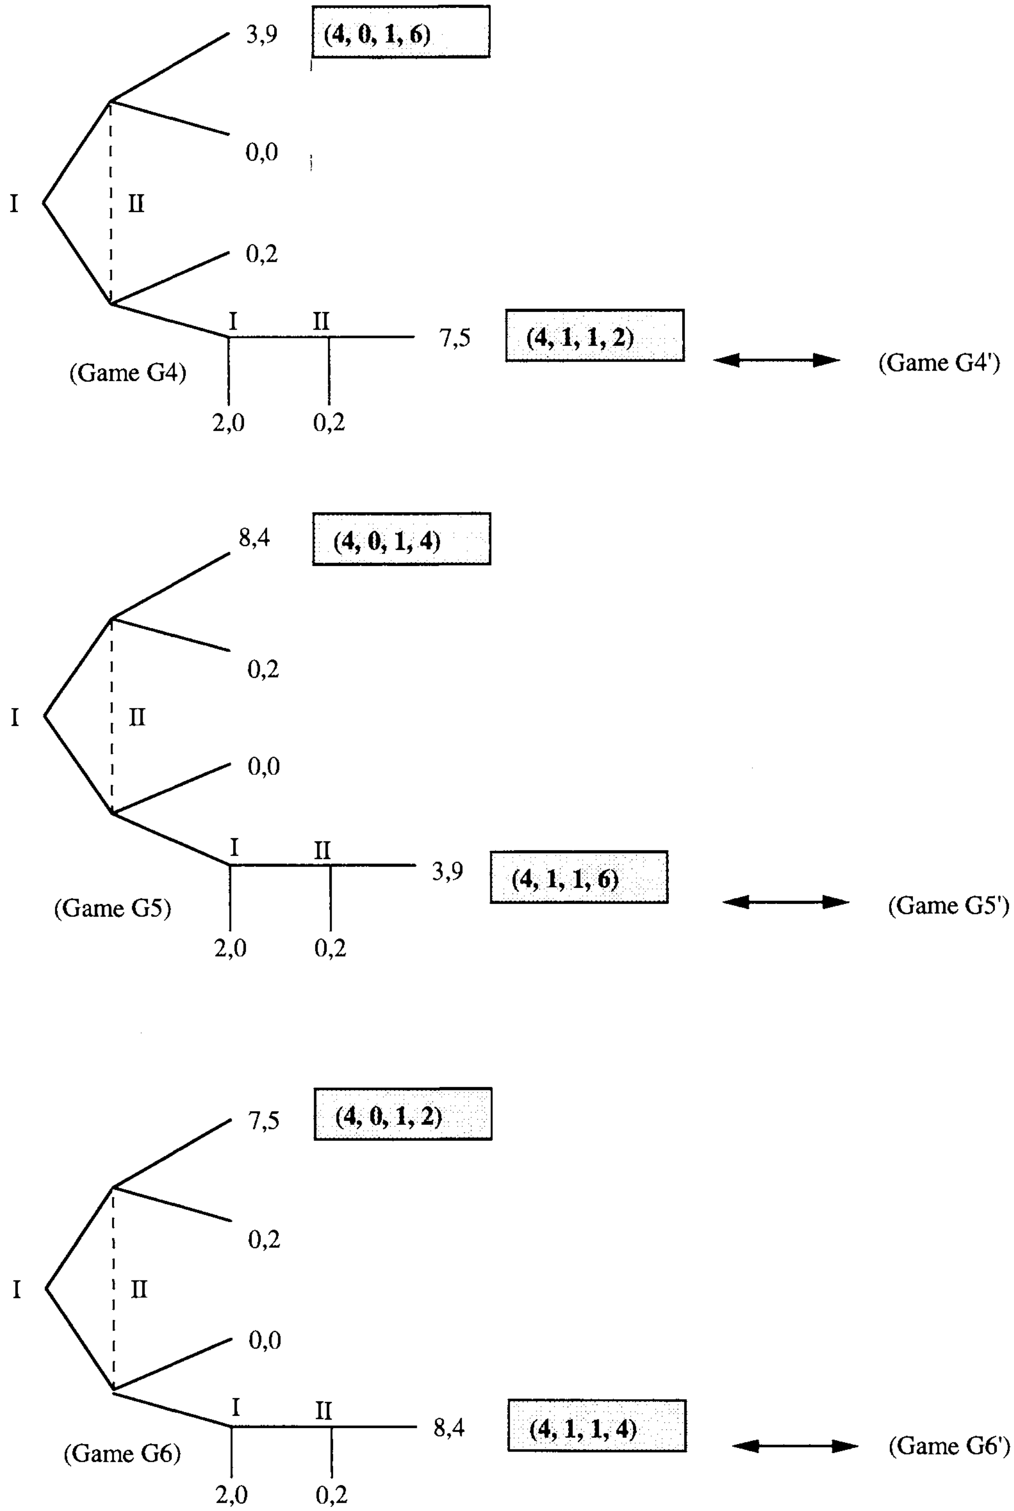

Figure 6 **Experimental Design for Determining the Effect of Iterative Knowledge of Strategies**

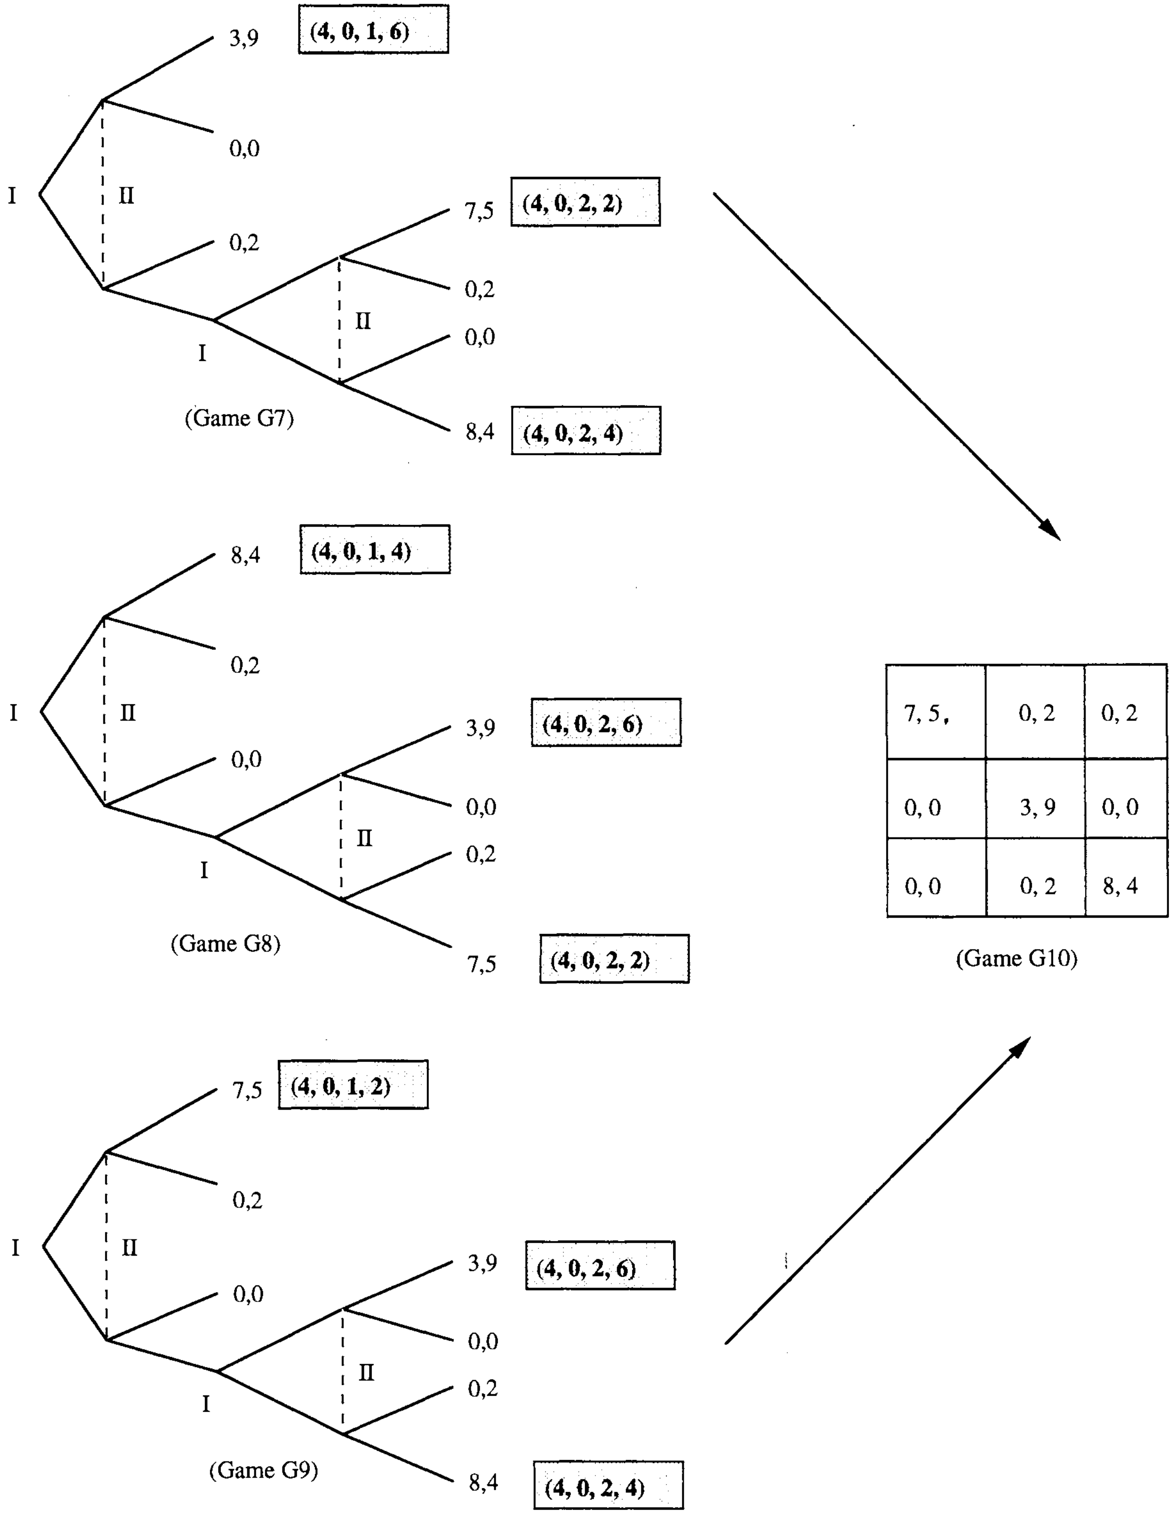

Table 1a **Frequency of Outcomes in Games G_1–G_3, G_1'**

Game Outcome						Round						
	1	2	3	4	5	6	7	8	9	10	11	Total
G_1												
SS	7	6	7	9	8	10	10	10	10	9	10	96
SC	1	2	0	1	1	0	0	0	0	0	0	5
CS	1	2	3	0	1	0	0	0	0	1	0	8
CC	1	0	0	0	0	0	0	0	0	0	0	1
G_1'												
SS	7	8	6	6	7	7	8	7	6	7	7	76
SC	4	3	4	3	4	0	1	3	2	1	2	27
CS	2	2	2	3	1	3	1	0	2	2	1	19
CC	2	2	3	3	3	5	5	5	5	5	5	43
G_2												
SS	9	9	10	14	14	14	14	15	15	15	15	144
SC	3	1	2	0	0	0	0	0	0	0	0	6
CS	3	5	3	1	1	1	1	0	0	0	0	15
CC	1	0	0	0	0	0	0	0	0	0	0	1
G_3												
SS	3	4	6	7	10	7	9	11	8	12	12	89
SC	3	2	4	2	1	5	5	1	2	1	2	28
CS	6	8	2	4	3	2	1	3	5	2	1	37
CC	3	1	3	2	1	1	0	0	0	0	0	11
Pooled												
SS	26	27	29	36	39	38	41	43	39	43	44	405
SC	11	8	10	6	6	5	6	4	4	2	4	66
CS	12	17	10	8	6	6	3	3	7	5	2	79
CC	7	3	6	5	4	6	5	5	5	5	5	56

subjects would be assigned a new pair member in each round. And they would never be paired with another subject more than once during the 10 rounds.

The role of each player (i.e., player 1 or 2) was randomly assigned and could change between rounds. Players were informed of their role for the round by looking at their worksheets; it was designated by the label "YOU," and their payoffs were in bold type.[6] Subjects selected strategies simultaneously. Administrators recorded the subject's choice, and that of the pair member, and informed the two subjects of their payoffs. Subjects received this information at the end of each round, before making next round's decision. Strategy choices and payoffs of other subject-pairs were not revealed to them. In multi-stage problems, if a subject chose a Stage I strategy that might result in the game proceeding to Stage II (depending on their pair member choice), they were asked to simultaneously reveal their Stage II strategy. We did this for two reasons: One, to speed up the experiment, (subjects get bored if there is a large time lapse between stages). Two, we could record the contingent strategies of subjects: Even if a subject's pair member did not choose a strategy that led to Stage II,

[6] Previous experiments show that subjects learn faster if they are assigned both roles (Binmore et al. 1985). We speculate this occurs because it makes it easier for them to "view" the game from the other's perspective.

Table 1b Frequency of Outcomes in Games G_4–G_6, G'_4–G'_6

Game					Round						
Outcome	1	2	3	4	5	6	7	8	9	10	11	Total
G_4												
SS	6	1	5	5	6	5	8	9	10	12	11	78
SC	2	8	0	2	1	3	3	2	2	2	1	26
CS	5	5	6	5	7	3	2	2	3	2	4	44
CC	9	7	10	10	11	8	10	10	9	8	8	100
Others	3	4	4	3	0	6	2	2	1	1	1	27
G'_4												
SS	7	8	9	6	5	6	7	6	7	8	7	76
SC	3	3	3	5	2	4	2	6	2	4	1	35
CS	3	1	1	2	5	2	1	1	3	1	3	23
CC	2	3	2	2	3	3	5	2	3	2	4	31
Others	0	0	0	0	0	0	0	0	0	0	0	0
G_5												
SS	14	13	13	15	17	17	18	18	19	18	18	180
SC	4	3	3	3	1	0	0	1	0	1	1	17
CS	0	3	3	2	1	0	1	1	1	0	0	12
CC	2	0	0	0	0	0	0	0	0	0	0	2
Others	0	1	1	0	1	3	1	0	0	1	1	9
G'_5												
SS	1	3	3	6	5	9	8	13	10	9	10	77
SC	5	5	1	1	1	1	4	0	0	1	0	19
CS	4	6	10	7	8	4	3	2	4	5	5	58
CC	5	1	1	1	1	1	0	0	1	0	0	11
Others	0	0	0	0	0	0	0	0	0	0	0	0
G_6												
SS	10	11	10	11	14	13	13	16	15	16	15	144
SC	5	5	2	2	2	2	4	2	4	2	1	31
CS	0	1	2	5	2	2	0	1	1	2	2	18
CC	0	0	1	0	0	1	0	0	0	0	1	3
Others	5	3	5	2	2	2	3	1	0	0	1	24
G'_6												
SS	8	8	10	7	9	9	11	9	12	12	11	106
SC	3	1	3	2	1	4	3	5	0	1	2	25
CS	3	5	0	3	5	1	0	1	2	1	2	23
CC	1	1	2	3	0	1	1	0	1	1	0	11
Others	0	0	0	0	0	0	0	0	0	0	0	0
Pooled												
SS	46	44	50	50	56	59	65	71	73	75	72	661
SC	22	25	12	15	8	14	16	16	8	11	6	153
CS	15	21	22	24	28	12	7	8	14	11	16	178
CC	19	12	16	16	15	14	16	12	14	11	13	158
Others	8	8	10	5	3	11	6	3	1	2	3	60

Table 1c Frequency of Outcomes in Games *G7–G9*

Game Outcome	Round 1	2	3	4	5	6	7	8	9	10	11	Total
G_7												
SS	8	7	9	11	9	10	9	13	11	12	9	108
SCt	5	4	6	6	6	5	7	5	7	6	9	66
SCb	3	5	3	1	5	2	1	2	4	1	2	29
CtS	3	1	2	2	0	2	1	0	1	3	1	16
CbS	0	1	2	0	0	0	1	1	0	1	1	7
CtCb	3	2	0	0	4	5	4	2	1	3	1	25
CbCt	3	1	0	4	2	2	3	2	1	0	2	20
CtCt	4	5	8	5	4	3	4	5	4	4	4	50
CbCb	1	4	0	1	0	1	0	0	1	0	1	9
G_8												
SS	3	5	6	7	10	10	13	14	13	15	13	109
SCt	0	2	1	0	1	0	0	0	0	0	0	4
SCb	5	5	6	3	3	2	3	1	0	1	1	30
CtS	3	2	1	1	3	3	0	0	3	1	2	19
CbS	4	4	3	3	1	4	0	2	1	2	2	26
CtCb	1	0	0	2	0	0	1	1	3	1	1	10
CbCt	1	0	0	0	0	0	1	1	0	0	0	3
CtCt	0	0	1	2	1	1	0	0	0	0	0	5
CbCb	3	2	2	2	1	0	2	1	0	0	1	14
G_9												
SS	11	11	11	7	9	10	9	16	14	16	21	135
SCt	3	1	2	4	3	0	4	2	1	0	0	20
SCb	2	3	3	3	3	3	1	2	1	2	1	24
CtS	9	10	9	10	4	9	10	5	9	7	3	85
CbS	1	1	1	1	4	2	3	2	1	2	0	18
CtCb	0	4	1	1	3	3	2	0	1	1	1	17
CbCt	0	0	0	0	1	2	0	0	0	0	0	3
CtCt	2	0	1	3	3	1	1	2	3	2	4	22
CbCb	2	0	2	1	0	0	0	1	0	0	0	6
Pooled												
SS	22	23	26	25	28	30	31	43	38	43	43	352
SCt	8	7	9	10	10	5	11	7	8	6	9	90
SCb	10	13	12	7	11	7	5	5	5	4	4	83
CtS	15	13	12	13	7	14	11	5	13	11	6	120
CbS	5	6	6	4	5	6	4	5	2	5	3	51
CtCb	4	6	1	3	7	8	7	3	5	5	3	52
CbCt	4	1	0	4	3	4	4	3	1	0	2	26
CtCt	6	5	10	10	8	5	5	7	7	6	8	77
CbCb	6	6	4	4	1	1	2	2	1	0	2	29

Table 2 **Chi-square Values and *t*-statistics for Hypotheses Testing**

Game	Hypothesis 1 Chi-Square	Hypothesis 2 *t*-statistics	Hypothesis 3 *t*-statistics	Hypothesis 3 95% Confidence Intervals	Game	Hypothesis 1 Chi-Square	Hypothesis 2 *t*-statistics	Hypothesis 3 *t*-statistics	Hypothesis 3 95% Confidence Intervals
G_1					G_5'				
1–6	92.3***	−8.4	8.4***	(0.79, 0.97)	1–6	23.2***	−2.5	2.5**	(0.52, 0.67)
7–11	142.2***	−9.8	9.6***	(0.89, 1.00)	7–11	79.0***	−8.0	7.5***	(0.75, 0.91)
1–11	228.4***	−12.8	12.7***	(0.87, 0.99)	1–11	72.0***	−4.8	5.7***	(0.65, 0.75)
G_1'					G_6				
1–6	21.0***	3.4***	3.4***	(0.55, 0.70)	1–6	92.4***	9.4***	9.4***	(0.76, 0.91)
7–11	29.9***	1.6	1.6	(0.49, 0.65)	7–11	143.2***	10.7***	10.7***	(0.82, 0.96)
1–11	47.9***	3.6***	3.6***	(0.55, 0.65)	1–11	228.5***	14.2***	14.2***	(0.81, 0.91)
G_2					G_6'				
1–6	95.6***	−8.6	8.6***	(0.81, 0.95)	1–6	50.0***	−6.4	6.1***	(0.73, 0.89)
7–11	142.2***	−9.8	9.8***	(0.91, 1.00)	7–11	95.2***	−8.5	8.5***	(0.77, 0.93)
1–11	230.5***	−12.9	12.9***	(0.88, 0.98)	1–11	130.3***	−10.4	10.5***	(0.77, 0.88)
G_3					G_7				
1–6	16.8***	−3.9	3.9***	(0.57, 0.72)	1–6	102.0***	0.7	5.8***	(0.43, 0.53)
7–11	83.3***	−8.5	8.5***	(0.77, 0.93)	7–11	141.8***	0.1	7.6***	(0.49, 0.59)
1–11	82.2***	−8.6	8.6***	(0.68, 0.79)	1–11	236.8***	0.6	9.4***	(0.47, 0.54)
G_4					G_8				
1–6	24.0***	3.4***	−3.3	(0.34, 0.46)	1–6	85.3***	−4.7	8.5***	(0.53, 0.65)
7–11	42.4***	−0.7	6.5***	(0.46, 0.59)	7–11	323.6***	−5.5	13.3***	(0.71, 0.84)
1–11	50.9***	2.0*	−2.0	(0.42, 0.50)	1–11	311.8***	−7.1	15.2***	(0.63, 0.72)
G_4'					G_9				
1–6	21.2***	3.9***	3.9***	(0.57, 0.72)	1–6	165.9***	10.0***	10.0***	(0.53, 0.63)
7–11	20.3***	3.1***	3.1***	(0.55, 0.71)	7–11	287.0***	13.2***	13.2***	(0.64, 0.75)
1–11	42.3***	5.0***	5.0***	(0.58, 0.69)	1–11	422.6***	16.3***	16.3***	(0.60, 0.67)
G_5									
1–6	162.5***	11.5***	11.5***	(0.82, 0.95)					
7–11	238.0***	13.1***	13.1***	(0.90, 1.00)					
1–11	475.4***	17.3***	17.3***	(0.87, 0.97)					

* $p < 0.05$.
** $p < 0.01$.
*** $p < 0.005$.

we knew the subject's strategy in Stage II. Payoff points were worth \$.10. After 10 rounds subjects would sum their points and multiply this amount by .1. These were their dollar earnings for that decision problem.

3.2. Experimental Design

Seventeen experimental sessions ($N = 187$) investigated the impact of task complexity and payoff disparity on subject behavior. Subjects usually participated in 3

Figure 7 Proportion of Simple Equilibrium over Time

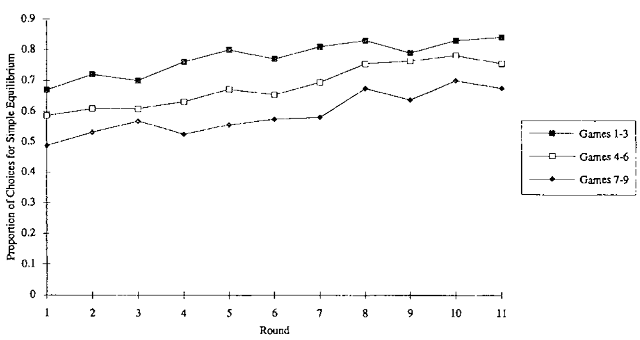

different games in each session. We use a Latin square design to control for possible order effects.[7]

Figures 4–6 show the games played by subjects. Games G_n and G'_n ($n = 1, 4, 5, 6$) are identical except the positions of the equilibria are switched. For example, game G'_4 is the same as G_4 except payoffs (7, 5) and (3, 9) are interchanged. Each equilibrium x is represented as a 4-dimensional vector (c_x, r_x, s_x, p_x). We did not want to confound results with simultaneous changes in task complexity metrics, so in each game only 1 task complexity metric, and payoff disparity were manipulated. c_x is manipulated in games G_1–G_3 and G'_1; r_x in games G_4–G_6 and G'_4–G'_6; and s_x in games G_7–G_9. For each equilibrium x, the shaded box gives the levels of task complexity and payoff disparity. For example, in Figure 4, game G_1, the upper equilibrium has a c index of 16, a r index of 0, a s index of 2, and a p index of 2.

4. Results

The frequency of outcomes across rounds is shown in Tables 1a–1c. Data are shown for individual games and are pooled across games which manipulate the same task complexity metric. For example, in Table 1c, games G_7–G_9 only manipulate s_x. The equilibrium with the smaller complexity metric is always designated as S (Simple), and the equilibria with the higher metric as C (Complex). So if both players 1 and 2 choose an action leading to the simple equilibrium, the outcome is labeled as SS. Similarly, if player 1 chooses an action leading to the complex equilibrium, but player 2 chooses an action leading to the simple one, the outcome is labeled as CS. For example, in round 1 of game G_5 (Table 1b), 14 subject pairs coordinated on the simple equilibrium; in 4 pairs, player 1 attempted to reach the simple equilibrium, while player 2 attempted to reach the complex one; and 2 pairs coordinated on the complex equilibrium. In games G_7–G_9, there are two complex equilibrium. The top complex equilibrium is labeled as C_t, the bottom as C_b.[8]

We test several hypotheses. Since all equilibria in our experimental games are subgame perfect, our first hypothesis examines whether the frequency distribution across outcomes is approximately uniform.

HYPOTHESIS 1. *The frequency distribution across outcomes is uniform.*

Casual observation of the "Total" column of Tables 1a–c, suggests that in all games, subjects have a systematic preference for one equilibrium: Subject-pairs strongly prefer the simple equilibrium outcome. In the pooled data, the simple equilibrium outcome (SS) is chosen at least 3 times more frequently than any other outcome. Also, in the pooled data across games with only 2 equilibria, the frequency of SS is greater than the sum of frequencies across all other outcomes.

Following normal procedure, we conduct chi-square tests on three different "cuts" of the data: the summed frequency of choices over the first 6 rounds, over the last 5 rounds, and over the entire 11 rounds. We did this to detect any changes in choice patterns over time. These goodness-of-fit tests examined whether the observed frequency of choices are consistent with the theoretical prediction.[9] We report the chi-square values in

[7] Unfortunately, in a few experimental sessions, subjects were not able to stay for the complete sequence of games. Chi-square tests suggest order effects in G_4, G_7, and G_9. Order did not appear to affect learning.

[8] In two experimental sessions involving G_4 and G_6, a few subjects chose nonequilibrium outcomes. We have no explanation for their choices. Apparently, this behavior caused other subjects to "retaliate" and adopt similar behavior. As shown in Table 1b, this behavior was primarily confined to early rounds and only accounted for 5% of the total choices.

[9] In games G_1–G_6, G'_1, G'_4–G'_6, if subjects were to randomize between equilibria, each equilibrium would be selected $0.25*N$ times, where N is the total number of subject pairs in the round. (In games G_4–G_6 and G'_4–G'_6, we assume that subjects will not choose a dominated strategy.) In games G_7–G_9, each equilibrium should be chosen $0.11*N$ times.

Table 3 **An Illustrative Example** ($\alpha = -0.2$, $\beta = 0.5$)

	Round (j)											
	0	1	2	3	4	5	6	7	8	9	10	11
Subject (i)	–	C	S	C	S	S	C	S	S	S	S	–
Pair Member	–	C	S	S	S	C	S	S	C	S	S	–
$\delta_i^j(S)$ see Eq. (4.1)	0	0	1	1	1	0	1	1	0	1	1	0
$\delta_i^j(C)$ see Eq. (4.1)	0	1	0	0	0	1	0	0	1	0	0	0
$n_i^j(S)$ see Eq. (4.2)	0	0	1	2	3	3	4	5	5	6	7	7
$n_i^j(C)$ see Eq. (4.2)	0	1	1	1	1	2	2	2	3	3	3	3
$m_i^j(S)$ see Eq. (4.3)	0	0	1	1	2	1	1	2	1	2	3	3
$m_i^j(C)$ see Eq. (4.3)	0	1	1	0	0	0	−1	−1	−1	−1	−1	−1
$p_i^j(S)$ (Cournot)	–	0.55	0.31	0.77	0.77	0.77	0.31	0.77	0.77	0.31	0.77	0.77
$p_i^j(S)$ (Fict. Play)	–	0.55	0.43	0.55	0.67	0.77	0.67	0.77	0.85	0.77	0.85	0.90
$p_i^j(S)$ (Pay. Reinf.)	–	0.55	0.43	0.55	0.67	0.77	0.67	0.77	0.85	0.77	0.85	0.90

Table 2. As shown in Table 2, in all games, and all three cuts of the data, subjects did not randomize across equilibria ($p < 0.001$); they appeared to use choice rules that led to the least task complex equilibrium.

However, subjects' choices across rounds are probably not independent. This dependency may overstate the chi-square values (and the t-statistics in Hypotheses 2 and 3). To address this issue, one must examine learning behavior across rounds, which most previous studies ignored. We use a simple maximum likelihood test to analyze learning behavior below. Still, we report these traditional chi-square and t-tests as benchmarks.

We examine several hypotheses below that may help explain subjects' behavior. As indicated above, some evidence suggests that players care about payoff disparity. The next hypothesis examines whether subjects consider payoff disparity in choosing actions:

HYPOTHESIS 2. *The proportion of subjects choosing the equilibrium with the lowest payoff disparity is significantly larger than 1/number of equilibria.*

A conservative test of whether subjects consider payoff disparity is to test whether the equilibrium with the lowest payoff disparity was chosen more often than the other equilibria. In games G_1–G_6, G_1', G_4'–G_6', there are 2 equilibria so we test whether the lowest payoff disparity equilibrium was chosen more than 50% of the time. Similarly, we test whether the low payoff disparity equilibrium was selected more than 33% of the time in games G_7–G_9. Table 2 shows the t-statistics. In games G_1', G_4, G_4', G_5, G_6, and G_9, subjects chose the low payoff disparity equilibrium significantly more often. In the remaining games (except G_7) they chose the low payoff disparity equilibrium significantly less often. Note that in all the games except G_4, in which subjects chose the low payoff disparity equilibrium significantly more often, this equilibrium was also the least task complex. In the last 5 rounds of game G_4, subjects did not choose the low payoff disparity equilibrium significantly more often. We can untangle these confounding effects by comparing the choice patterns in game pairs G_1–G_1', G_4–G_4', G_5–G_5', and G_6–G_6'. In 3 of the 4 cases, subjects overwhelmingly chose the least task complex equilibrium whether it was high or low payoff disparity. In the game pair G_4–G_4', subjects switched toward the least task complex equilibrium, but not overwhelmingly. Overall, these results reject Hypothesis 2.

The above discussion suggests that subjects may adopt effort conserving rules that lead them to the equilibrium with the lowest task complexity. We formally test this in Hypothesis 3.

HYPOTHESIS 3. *Subjects will adopt effort conserving rules that lead to the less task complex equilibrium.*

Using a t-test, we test whether the low task complexity equilibrium was chosen more than 1/number of equilibria times (50% in games G_1–G_6, G_1', G_4'–G_6' and

Table 4 Log Likelihoods and Parameter Estimates for No Learning and the Three Learning Models

	Uniform Prior & No Learning	Biased Prior & No Learning			Uniform Prior with Cournot			Biased Prior with Cournot				
Game	Ln(L)	α^*	Ln(L)	χ^2	β^*	Ln(L)	χ^2	α^*	β^*	Ln(L)	χ^2_α	χ^2_β
G_1	−152.5	−2.6	−54.8	195.4*	2.40	−79.36	146.4*	−2.0	1.18	−49.84	59.0*	9.9*
G_1'	−228.7	−0.4	−222.1	13.2*	1.15	−190.61	76.18*	−0.3	1.10	−187.75	5.7*	68.7*
G_2	−228.7	−2.6	−80.8	295.8*	2.20	−129.36	198.68*	−2.1	0.98	−76.13	106.5*	9.3*
G_3	−228.7	−1.0	−190.4	76.6*	0.83	−207.36	42.68*	−0.9	0.53	−184.35	46.0*	12.1*
G_4	−381.2	0.3	−376.0	20.4*	1.15	−317.30	127.8*	0.2	1.12	−315.05	4.5	121.9*
G_4'	−228.7	−0.6	−215.2	27.0*	0.80	−208.19	41.02*	−0.5	0.73	−200.94	14.5*	28.5*
G_5	−305.0	−2.2	−140.8	328.4*	1.60	−217.66	174.68*	−2.0	0.37	−139.59	156.1*	2.4
G_5'	−228.7	−0.8	−201.7	54.0*	0.73	−212.03	33.34*	−0.7	0.53	−194.86	34.3*	13.7*
G_6	−305.0	−1.4	−218.8	172.4*	1.08	−258.99	92.02*	−1.2	0.52	−212.57	92.8*	12.5*
G_6'	−228.7	−0.8	−171.8	113.8*	1.17	−189.45	78.5*	−1.0	0.75	−161.33	56.2*	20.9*
G_7	−725.1	−0.7	−683.4	83.4*	0.93	−669.15	111.9*	−0.6	0.80	−644.22	49.9*	78.2*
G_8	−483.4	−1.5	−369.4	228.0*	1.33	−408.31	150.18*	−1.2	0.85	−348.08	120.1*	42.6*
G_9	−725.1	−1.2	−604.0	242.2*	0.98	−662.55	125.1*	−1.1	0.57	−586.76	151.6*	34.48*

33% in games G_7–G_9). Table 2 shows the t-statistics; except for G_4, the values are significant at the .01 level. Even in G_4 the t-statistic is significant for the last 5 rounds. Also, in all games, except G_1', the t-statistics increase in rounds 7–11, relative to rounds 1–6. This suggests that in rounds 7–11 subjects chose the less task complexity equilibrium more often relative to behavior in rounds 1–6.

We then construct 95% confidence intervals for the probability of choosing the lowest task complexity equilibrium. In all games except G_4, the lower bounds of the confidence intervals are greater than 1/Number of equilibrium. In 9 out of 13 games, the lower bound is at least 0.64. This suggests subjects are adopting effort conserving rules that lead to the least task complexity equilibrium.

Figure 7 shows that the proportion of subjects choosing the least task complex equilibrium increases over time. We conduct chi-square tests to determine whether the choice pattern of outcomes of subject-pairs in the last 5 rounds was significantly different from that in the first 6 rounds (as evidenced by confidence intervals shown in Table 2). The chi-square values are 23.5 ($p < 0.01$) for games G_1–G_3 and G_1'; 43.6 ($p < 0.01$) for games G_4–G_6 and G_4'–G_6'; and 23.0 ($p < 0.01$) for games G_7–G_9. This result indicates that subjects did choose the least task complex equilibrium more often in later rounds. Overall, results do support Hypothesis 3. Subjects do appear to use task complexity as a selection criterion in our coordination games.

4.1. Game-form Independence

We also test whether choice behavior is game form dependent. As indicated above, a reduced normal-form game can have multiple extensive-form equivalents. Games G_7 and G_9 have an identical reduced normal-form game (G_{10}). At the end of some experimental sessions, we presented subjects with game G_{10}. Observed results show the majority of subjects chose the equilibrium (3, 9). We performed a chi-square test for the pooled results of rounds 1–6, 7–11, and 1–11. The test shows subject behavior is significantly different at the .001 level in game pairs G_7–G_9 (rounds 1–6—47.79; 7–11—105.7; 1–11—145.142) and G_9–G_{10} (1–6—29.41; 7–11—112.2; 1–11—126.3). In the strategically equivalent games G_7 and G_{10}, behavior is similar (1–6—2.8; 7–11—5.1; 1–11—3.6). However, as discussed below, subjects may apply different equilibrium selection criteria in the two games. This result and previous experiments (Schotter, Weigelt and Wilson 1994) suggest that normal-form based equilibrium concepts (e.g.,

Table 4 *Continued*

Uniform Prior with Fictitious Play			Biased Prior with Fictitious Play					Uniform Prior with Payoff Reinforcement			Biased Prior with Payoff Reinforcement				
β^*	Ln(L)	χ^2	α^*	β^*	Ln(L)	χ^2_α	χ^2_β	β^*	Ln(L)	χ^2	α^*	β^*	Ln(L)	χ^2_α	χ^2_β
0.88	−50.4	204.2*	−1.2	0.52	−44.5	11.8*	20.6*	0.88	−50.0	204.2*	−1.2	0.52	−44.5	11.0*	20.6*
0.40	−159.3	139.2*	0.0	0.40	−159.3	0.0	125.6*	0.40	−159.3	138.8*	0.0	0.40	−159.3	0.0	125.6*
0.90	−74.8	307.9*	−1.2	0.54	−65.7	36.4*	30.2*	0.90	−74.77	307.9*	−1.2	0.54	−65.7	18.1*	30.2*
0.42	−181.0	95.4*	−0.6	0.31	−173.0	16.0*	34.8*	0.42	−181.0	95.4*	−0.6	0.31	−173.0	16.0*	17.4*
0.53	−249.3	263.8*	0.0	0.53	−249.3	0.0	253.4*	0.53	−249.3	263.8*	0.0	0.53	−249.3	0.0	253.4*
0.38	−180.2	97.0*	−0.2	0.35	−179.5	1.4	71.4*	0.38	−180.2	97.0*	−0.2	0.35	−179.5	1.4	71.4*
0.61	−145.6	318.8*	−1.3	0.29	−129.7	31.8*	22.2*	0.61	−145.6	318.8*	−1.3	0.29	−129.7	31.8*	22.2*
0.25	−210.5	36.4*	−0.7	0.13	−197.8	25.4*	7.8*	0.25	−210.5	36.4*	−0.7	0.13	−197.8	25.4*	7.8*
0.46	−210.7	188.6*	−0.7	0.29	−201.9	17.6*	33.8*	0.46	−210.7	188.6*	−0.7	0.29	−201.9	17.6*	33.8*
0.54	−143.0	171.4*	−0.5	0.44	−138.6	8.8*	66.4*	0.54	−143.0	171.4*	−5.0	0.44	−138.6	8.8*	66.4*
0.42	−587.7	274.8*	−0.3	0.40	−580.7	14.0*	205.4*	0.29	−672.3	105.6*	−0.5	0.23	−655.6	33.4*	55.6*
0.52	−352.1	262.6*	−0.8	0.37	−328.3	47.6*	82.2*	0.35	−427.4	112.0*	−1.3	0.18	−357.9	139.0*	23.0*
0.50	−559.4	331.4*	−0.5	0.39	−546.2	26.4*	115.6*	0.31	−672.2	105.8*	−1.1	0.13	−596.4	151.6*	15.2*

proper, persistent) probably do not apply to extensive form games.

4.2. Simple Learning Rules and Choice Behavior

As shown in Figure 7, subject behavior is path dependent: the frequency of choosing the least task complex equilibrium in later rounds appears dependent on its choice frequency in earlier rounds. This illustrates the fore-mentioned dependency issue, and suggests subjects adjust their behavior across rounds. To gain an understanding of these behavioral dynamics, we fit three simple learning models to subjects' behavior. These learning models suggest ways subjects may update the probability of choosing the least task complex equilibrium based on history of play. It also relaxes the independence assumption of behavior across rounds. Two learning models are suggested by economists, one by psychologists, and all are embedded in a multi-logit model (Cox 1958).

Before we describe the model, we define some notation. Our experimental games have either 2 (G_1 to G_6) or 3 (G_7 to G_9) equilibria. For expository purposes, we consider games with 2 equilibria. They are labeled as S (simple or the least task complexity equilibrium) or C. At round j ($j = 1, \ldots, 11$), subject i ($i = 1, \ldots, 11$) will choose equilibrium x with a probability $p_i^j(x)$. Obviously, $\Sigma_{x=S,C}\, p_i^j(x) = 1$.

4.3. Cournot

The Cournot model suggests behavior in the current round is solely dependent on behavior in the immediate previous round (Cournot 1838). Specifically, subjects assume their current and immediate previous round pair members will behave identically. Accordingly, subjects choose a best response based on this conjecture.

Let $\delta_i^j(x) = 1$ if subject i's opponent chooses x ($x = S$, C) in round j and 0 otherwise. If the subject follows Cournot dynamics, she will choose equilibrium S in round $j + 1$ with a probability given by.[10]

$$p_i^{j+1}(S) = \frac{e^{\beta\delta_i^j(S)}}{e^{\beta\delta_i^j(S)} + e^{\alpha+\beta\delta_i^j(C)}} \tag{4.1}$$

where $p_i^{j+1}(C) = 1 - p_i^{j+1}(S)$. α represents a subject's relative inclination for a more task complex equilibrium.

[10] Games G_7–G_9 have three equilibria. Here,

$$p_i^{j+1}(S) = \frac{e^{\beta\delta_i^j(S)}}{e^{\beta\delta_i^j(S)} + e^{\alpha+\beta\delta_i^j(C_1)} + e^{\alpha+\beta\delta_i^j(C_2)}},$$

$$p_i^{j+1}(C_1) = \frac{e^{\alpha+\beta\delta_i^j(C_1)}}{e^{\beta\delta_i^j(S)} + e^{\alpha+\beta\delta_i^j(C_1)} + e^{\alpha+\beta\delta_i^j(C_2)}},$$

and $p_i^{j+1}(C_2) = 1 - p_i^{j+1}(S) - p_i^{j+1}(C_1)$. We assume both complex equilibria have the same α.

Figure 8 **Invariant Transformations on Extensive-form Games**

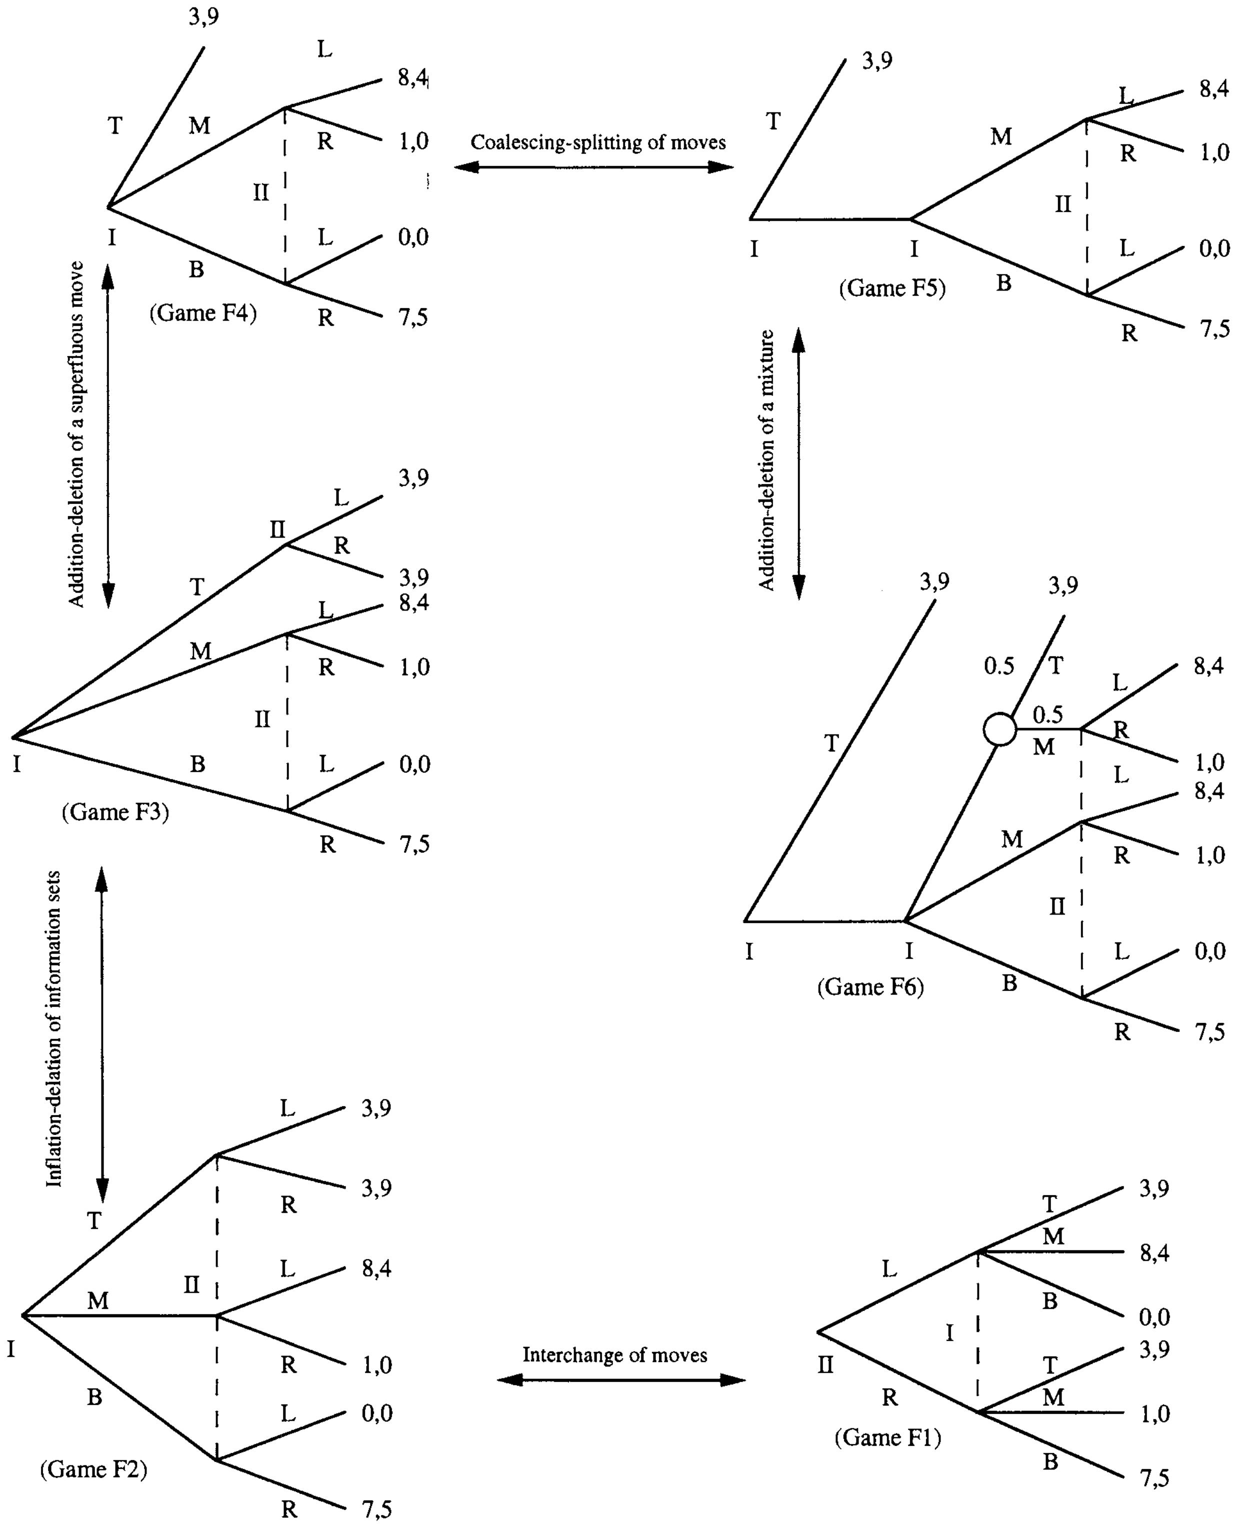

If a subject is averse to task complexity, $\alpha < 0$ and if she prefers task complexity, $\alpha > 0$. $\beta > 0$ captures the degree of path dependency; a higher β implies greater path-dependency. If subject i's pair member chooses s in round j, then i will choose s in round $j + 1$ with a probability $e^{\beta}/(e^{\beta} + e^{\alpha})$. As $\beta \to \infty$, then $p_i^{j+1}(S) \to 1$.

4.4. Fictitious Play

The Fictitious Play model is suggested by Brown (1951). This model assumes a subject's behavior in the current round is only dependent on the cumulative frequencies of choices across their previous pair members. The higher a choice's cumulative frequency, the higher the probability subject i will play that choice to coordinate with her pair member. Indeed, if a subject follows Fictitious Play, she weighs every round of the history equally.

Let subject i's pair members choose equilibrium x with $n_i^j(x)$ times by the end of round j. Subject i will choose equilibrium S with a probability given by:

$$p_i^{j+1}(S) = \frac{e^{\beta n_i^j(S)}}{\sum e^{\beta n_i^j(S)} + e^{\alpha + \beta n_i^j(C)}} \tag{4.2}$$

where $p_i^{j+1}(C) = 1 - p_i^{j+1}(S)$. For instance, at the end of round 4, if subject i's pair members have chosen S three times and C once, then i will choose S in round 5 with a probability $e^{3\beta}/(e^{3\beta} + e^{\alpha+\beta})$. As $\beta \to \infty$, then $p_i^{j+1}(S) \to 1$.

4.5. Payoff Reinforcement

The previous two models suggest subjects adjust their choices based on choices of pair members. The Payoff Reinforcement model, as suggested by Bush and Mosteller (1955), assumes subjects adjust their choices dependent on the reinforcement they receive (i.e., monetary payoffs).

Let subject i choose equilibrium x in round j. She can receive either a positive (i.e., her pair member chooses x) or a negative (i.e., her pair member chooses not x) reinforcement. Let the net reinforcement for choosing equilibrium x up to and including round j be $m_i^j(x)$. Subject i in round $j + 1$ will choose equilibrium S with the following probability:

$$p_i^{j+1}(S) = \frac{e^{\beta m_i^j(S)}}{\sum e^{\beta m_i^j(S)} + e^{\alpha + \beta m_i^j(C)}} \tag{4.3}$$

Hence, as the net reinforcement of an equilibrium increases, the probability of its choice increases. For instance, at the end of round 4, if subject i has chosen S three times and it is always positively reinforcement, C once and it is negatively reinforced, then i will choose S in round 5 with a probability $e^{3\beta}/(e^{3\beta} + e^{\alpha-\beta})$. As $\beta \to \infty$, then $p_i^{j+1}(S) \to 1$.

4.6. Maximum-likelihood Estimation

Let the likelihood function be $L(\alpha, \beta)$, which is the product of the probabilities of all subjects across all rounds. Let $\pi_i^j(x) = 1$ if subject i chooses x in period j and $\pi_i^j(x) = 0$ otherwise. Then $L(\alpha, \beta)$ is:

$$L(\alpha, \beta) = \prod_{i=1}^{11} \left[\prod_{j=1}^{11} (\pi_i^j(S) p_i^j(S) + \pi_i^j(C) p_i^j(C)) \right]. \tag{4.4}$$

Different learning models generate a different $p_i^j(x)$ dependent on past plays. Table 3 gives a numerical example.[11] Note that in all three models, subjects' behavior are path dependent because $p_i^j(x)$ depends on the history of plays. We determine α and β to maximize the likelihood function. The model with full specification including both α and β nests three simpler models ($L(0, 0)$, $L(\alpha, 0)$, $L(0, \beta)$). We test below whether each of these nested models is statistically rejected or not.

Table 4 presents the maximum log-likelihoods under the different learning models. Panels 1–2 show the log-likelihoods across games when there is no path-dependency or learning. In Panel 1, subjects are assumed to have no "biased prior" between equilibria (i.e., Log($L(0, 0)$)) and in Panel 2 they do have a "biased prior" (i.e., Log($L(\alpha, 0)$)). Intuitively, subjects will have a biased prior if they have a stronger tendency towards a particular equilibrium. Panels 3 and 4 show the log-likelihoods under Cournot model without and with "biased prior" respectively (Log($L(0, \beta)$), Log($L(\alpha, \beta)$)). Panels 5–6 and 7–8 are identical to Panels 3–4 except they assume different learning dynamics (5–6 for Fictitious Play, and 7–8 for Payoff Reinforcement). Highlights of this table include:

[11] It is worth noting that while the Payoff Reinforcement model is cognitively more complex than the Fictitious Play model, their predictions are identical in two-equilibrium games regardless of history of plays. In three-equilibrium games, their predictions are different.

• Comparing the maximum log-likelihoods which result from either including a "biased prior" or not, we clearly see the maximum log-likelihood is significantly greater when there is a "biased prior." This is true across all learning models. Also, the α^j estimates are nonpositive in 50 out of 52 cases (i.e., subjects are task complexity averse). Thus, after adjusting for dependency of choices across rounds, our result still rejects Hypotheses 1 and 2 and supports Hypothesis 3.

• Comparing the maximum log-likelihoods in Panels 4, 6, and 8 with those in Panel 2, we clearly see subjects' choices are path dependent. In every single case, [$\beta > 0$ and significant.]

• In games with two equilibria (G_1 to G_6), the Fictitious Play and Payoff Reinforcement models predict significantly better than Cournot. In games with three equilibria (G_7 to G_9), the Fictitious Play model predicts better than Cournot which in turn predicts better than Payoff Reinforcement.

5. Discussion

Our task complexity measures assume that when faced with multi-stage games, decision-makers decompose games into stages. This interpretation of the decision-making process inherently assumes players are boundedly rational (i.e., procedurally rational, see Simon 1982) and may use effort conserving rules in their information processing. The standard game theory assumption of collapsing game stages into a single strategy choice (i.e., the transformation of extensive-form to normal-form) assumes all players are infinitely sophisticated along our three dimensions of task complexity. This implies players are indifferent to task complexity.

We conducted experiments to test whether an equilibrium's task complexity influences its selection. Results suggest it does; subjects are averse to task complexity. This aversion to task complexity is analogous to commonly-found risk averse behavior. Previous work has clearly shown that individuals consider risk in making choices between gambles. Economists capture this choice selection criteria via the utility function. By analogy, we speculate a plausible way to model bounded rationality is to capture an equilibrium's task complexity via a value function. We make an initial attempt to measure task complexity using three simple metrics.

The manipulation of the task complexity measures in our experiments was straightforward; for the c measure we added 2 more alternatives, and for the r and s measures we added 1 more level of iterative mutual knowledge of rationality and strategies. Even with this simple manipulation, subjects chose the least task complex equilibrium significantly more often. This result suggests our subjects may have either high costs of information processing or limited short-term memory for holding all the relevant information for executing a sophisticated rule of behavior (Miller 1956). Consequently, they are inclined to adopt simple, effort-conserving heuristics.

We also test whether subjects consider payoff disparity in equilibrium choice. Although results did not support Hypothesis 2, in two games (G_4 and G_7), subjects appear to choose the least payoff disparity equilibrium more often in early rounds.[12] But clearly, as a selection criterion it is significantly less important than task complexity. We conjecture subjects sometimes consider payoff disparity, and they make tradeoffs between task complexity and payoff disparity.

We focus on first-round behaviors only to understand how subjects make this tradeoff before learning takes place. That is, we set $\beta = 0$ and decompose α into

$$\alpha_c(c_C - c_S) + \alpha_r(r_C - c_S) + \alpha_s(s_C - s_S) + \alpha_p(p_C - p_S)$$

in Eq. (4.1)–(4.3). Subject i is predicted to choose the simple equilibrium in round 1 with a probability given by:

$$p_i^1(S) = \frac{1}{1 + e^{\alpha_c(c_C-c_S)+\alpha_r(r_C-r_S)+\alpha_s(s_C-s_S)+\alpha_p(p_C-p_S)}}. \quad (5.1)$$

If the subject is task complexity averse (loving) then α_c, α_r, $\alpha_s < 0 (> 0)$. If she dislikes (likes) payoff disparity, then $\alpha_p < 0 (\alpha_p > 0)$. If α_c, α_r, $\alpha_s = 0$, then the subject is indifferent to task complexity in choosing equilibria. If $\alpha_p = 0$, the subject does not care about the payoffs of others.[13] If the subject dislikes task complexity more than payoff disparity, α_c, α_r, and α_s will have higher absolute values than α_p.

[12] Note in both games, the simple equilibrium is (3, 9), which is also the equilibrium with the highest disparity in payoffs.

[13] We re-scaled the task complexity metrics to ensure that all variables have identical ranges.

We derive maximum likelihood estimates for α_c, α_r, α_s, and α_p using the proportion of equilibrium selection in the first round. In games G_1–G_3, G_1' our design allows us to estimate a pooled α_c and α_p; in games G_4–G_6 and G_4'–G_6', a pooled α_r and α_p; and in games G_7–G_9, a pooled α_s and α_p. The maximum likelihood estimates are $\alpha_c = -0.18$, $\alpha_p = 0.00$ (G_1–G_3); $\alpha_r = -0.18$, $\alpha_p = -0.14$ (G_4–G_6); and $\alpha_s = -0.17$, $\alpha_p = -0.14$ (G_7–G_9). Maximum likelihood ratio tests reveal that non-zero αs are significant at least at the .01 level. These result suggest subjects do consider payoff disparity but it is not as important as task complexity.

Results also indicate that subjects' behavior is gameform dependent. This result casts doubt on the descriptive validity of equilibrium concepts that are insensitive to game form. Experimenters and theorists should use caution in generalizing results across game forms.

Our results strongly suggest choices of subjects are path-dependent. This points out the danger of pooling data across rounds, although our results remain significant after accounting for learning. While we test only three learning rules, the simple methodology we use seems appropriate for other learning dynamics and experimental settings.

Also, the methodology reveals an interesting phenomenon about learning in simple coordination games. In games with two equilibria (G_1 to G_6), the more complex Payoff Reinforcement model fits the data better than the simple Cournot model whereas in games with three equilibria, the reverse is true (see Table 3). However, the moderately complex Fictitious Play model fits the best in all games. This finding suggests subjects did not use sophisticated learning rules in our experimental games. It will be worthwhile to investigate whether the same result would generalize to more complex games.

5.1. Harsayni-Selten 1-Point Solution Theory

Harsanyi and Selten (1988) develop a general theory for equilibrium selection based on the concepts of payoff and risk dominance and a mental reasoning (tracing) procedure. This theory is the only one that predicts a unique equilibrium in any finite game. Games G_1–G_9 have no payoff dominant equilibrium, but each game has an unique risk dominant one. For example in game G_8, equilibrium (3, 9) risk dominates (7, 5) in Stage II and it dominates (8, 4) in Stage I. The order of risk dominance between two equilibria can be defined using the evolutionary game theoretic concept of "resistance." A population of players can either play equilibrium E_1 or E_2. Imagine most players are selecting $E1$, except for a fraction of "mutant" players selecting E_2. The resistance of E_1 against E_2 is the maximum allowable fraction of mutants that can play E_2 without inducing present E_1 players to switch to E_2. So, E_1 risk dominates E_2, if and only if the resistance of E_1 against E_2 is greater than the resistance of E_2 against E_1 (see Myerson 1992, p. 118 for a more formal definition). Risk dominance predicts subjects will select the equilibrium with the highest resistance (3, 9—G_1–G_2, G_4–G_5, G_7–G_9; 8, 4—G_3 and G_6).

Surprisingly, in our experiments, this prediction is consistent with player 2 choosing her minimax strategy. In games G_4, G_7–G_9, there is an unique minimax strategy for player 2. For instance, in game G_7 (see Figure 6), player 2 can guarantee herself a payoff of 2 by choosing the upper branch (in Stage I). If Player 1 anticipates this, he will also choose the upper branch which will result in a payoff of (3, 9). In the remaining games, if subjects backwardly induct then player 2 has an unique minimax Stage I strategy. For example, in game G_5 (see Figure 5), player 2 can guarantee herself a payoff of 2 by choosing the bottom branch. If player 1 realizes this, he should choose the bottom branch to maximize his payoff.[14]

Table 1 shows that behavior in the last 5 rounds of games G_4 and G_7 is consistent with this interpretation. But, results from the remaining games are inconsistent with the risk dominance or minimaxing behavior. Hence, the results cast doubts on the descriptive validity of Harsanyi-Selten's 1 point solution theory.

5.2. Future Research

Our results strongly suggest the three task complexity metrics can predict choice behavior. An important extension of this work is to decompose our complexity constructs into more basic cognitive operations (see Johnson and Payne 1985). This will allow us to investigate how an increase in c, r, or s affect the number of basic operations.

As previously indicated, our task complexity metrics are sensitive to game form invariant transformations.

[14] Recall that subjects switched roles during the experiment.

Figure 8 shows 5 such game form transformations: an extension of our research will be to examine how these transformations affect subjects' equilibrium selection. Theoretically, all these transformations result in strategically equivalent or perfectly similar games. Our experiments suggest these transformations may result in games perceived to be dissimilar by subjects. Besides game form, games can be similar in other aspects; solution principles (e.g., iterative dominance, backward induction), nature of conflict (e.g., prisoner's dilemma, coordination games), order of payoffs. Like game-form dependency, we believe similarity is empirically important because it affects how individuals transfer learning across games. We plan to conduct a series of experiments to determine the major constituents of similarity as perceived by subjects.[15]

[15] This research was supported by the Sol C. Snider Entrepreneurial Center and the Reginald H. Jones Center for Management Policy, Strategy, and Organization. We thank Colin Camerer, Andy Schotter, Sushil Bikhchandani, the department editor, a reviewer, and seminar participants at University of Pennsylvania and Northwestern University for their helpful comments.

Appendix

Instructions

This is an experiment in decision-making. Several research foundations have provided funds for these experiments. If you follow the instructions, and make good decisions, you may earn a sizable amount of money. The amount of money you earn depends on your choices and the choices of other subjects in the experiment.

Your Decision Problem

During each round, you will be paired with a randomly selected subject. This subject will be your pair member for that round. In each round you and your pair member will be presented with the identical decision problem. Your role for that round is designated by the title of You. Your pair member's role is designated as "Pair member." Roles are randomly assigned, so your role may change from round-to-round.

In each decision round, you and your pair member will simultaneously choose between your action alternatives. Since you and your pair member have worked in privacy, after making your choice, an administrator will come around and record your choice. We will then match your choice with that of your pair member, and compute the point payoffs to all subjects. The administrator will then privately tell you what your point total is. You will record this on the appropriate line of you worksheet. We will then begin the next round.

The experiment will last for 10 rounds. In each round **you will be paired with a different pair member. That is, in this problem, another subject will never be your pair member more than once.** There are 11 subjects in the room. So each subject will 'sit out' one round. You will be informed of the round you will sit out by a blank sheet of paper in your folder. During this round you will not participate in the experiment.

The payoffs resulting from the choices of you and your pair member are described in the figures below. That is, the figure tells you how many points each of the possible decision choice pairs are worth to you. Note again that your payoffs are always in bold.

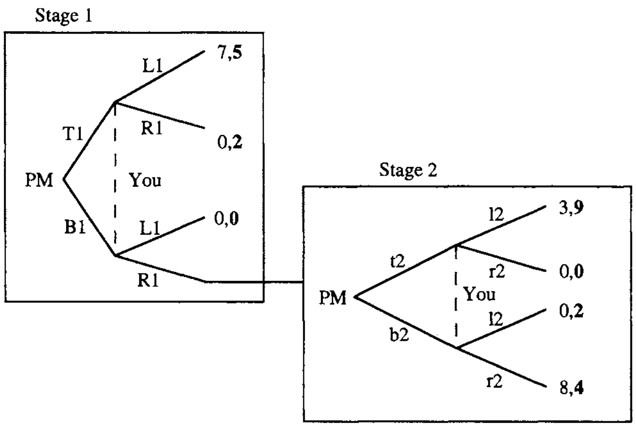

Some Examples

Suppose you choose R1 and your pair member chooses T1. Then you would receive 2 points and your pair member would receive 0 points.

Suppose you choose R1 and your pair member chooses B1. Because you chose R1, you would then have to make a choice in Stage 2. Suppose you choose r2 in Stage 2, and your pair member chooses b2. Then you would receive 4 points, and your pair member would receive 8 points.

After we have finished with round 1, we will proceed to round 2. The procedure is exactly the same except you will be randomly paired with another subject as your pair member, and you will be randomly assigned a role. Your role may change. You will not play this game against the same person more than one time.

Payoffs

Your dollar earnings for the experiment are determined as follows. First, we will sum up your point total over the 10 rounds. Then we will multiply this sum by $.10. We will then pay you this sum as you leave the experiment. Note that the more points you earn, the more money you will receive.

References

Banks, Jeffrey and J. Sobel, "Equilibrium Selection in Signaling Games," *Econometrica*, 55 (1987), 647–662.

Binmore, Kenneth, *Fun and Games*, D. C. Heath, Lexington, MA, 1992.

——, A. Shaked, and J. Sutton, "Testing Non-cooperative Bargaining Theory: A Preliminary Study," *American Economic Rev.*, 75 (1985), 1178–1180.

Brandenburger, Adam, "Knowledge and Equilibrium in Games," *J. Economic Perspectives*, 6 (1992), 83–101.

Brown, G., "Iterative Solution of Games by Fictitious Play," in *Activity Analysis of Production and Allocation*, Wiley, New York, 1951.

Bush, R. R. and F. Mosteller, *Stochastic Models for Learning*, Wiley, New York, 1955.

Cho, I. and David Kreps, "Signaling Games and Stable Equilibria," *Quarterly J. Economics*, 102 (1987), 179–221.

Cournot, A., *Recherches sur les principes mathematiques de la theorie des richesses*, translated into English by N. Bacon as *Researches in the Mathematical Principles of the Theory of Wealth*, Haffner, London, 1960.

Cox, D. R., "The Regression Analysis of Binary Sequences," *J. Royal Statistical Society*, 1958, 215–232.

Dalkey, Norman, "Equivalence of Information Patterns and Essentially Determinate Games," *Contributions to the Theory of Games II, Ann. Math. Studies* 28, Princeton University Press, Princeton, NJ 1953.

Fundenberg, Drew and Jean Tirole, *Game Theory*, MIT Press, Cambridge, MA, 1992.

Harsanyi, John and Reinhard Selten, *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge, MA, 1988.

Johnson, E. and J. Payne, "Effort and Accuracy in Choice," *Management Sci.*, 31, 4 (1985), 395–414.

Kagel, John and Alvin Roth (Eds.), *The Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ, 1995.

Kreps, D., "Out-of-equilibrium Beliefs and out-of-equilibrium Behavior," *The Economics of Missing Markets, Information, and Games*, Frank Hahn (Ed.), Clarendon Press, Oxford, 1990a.

——, *Game Theory and Economic Modelling*, Oxford University Press, Oxford, 1990b.

Lucas, Robert, "Adaptive Behavior in Economic Theory," *J. Business*, 59 (1986), s385–s400.

Marimon, R., E. McGrattan, and T. Sarget, "Money as a Medium of Exchange in an Economy with Artificial Intelligent Agents," *J. Economic Dynamics and Control*, 14 (1990), 329–374.

Maynard Smith, J., *Evolution and the Theory of Games*, Cambridge University Press, Cambridge, 1982.

Miller, G., "The Magical Number 7 Plus or Minus 2," *Psychological Rev.*, 63 (1956), 81–97.

Myerson, Roger, *Game Theory*, Harvard University Press, Boston, MA, 1992.

Rubinstein, Ariel, "Finite Automata Play in Repeated Prisoner's Dilemma," *J. Economic Theory*, 39 (1986), 83–96.

Schotter, Andrew, Keith Weigelt, and Charles Wilson, "A Laboratory Investigation of Multi-person Rationality and Presentation Effects," *Games and Economic Behavior*, 6 (1994), 445–468.

Selten, Reinhard, "Spieltheoretische behandlung eines oligopolmodells mit nachfragetragheit," *Zeitschrift fuer die gesampte Staatswisenschaft*, 121 (1965), 301–324, 667–689.

Simon, H., *Models of Bounded Rationality*, 2 vols., MIT Press, Cambridge, MA, 1982.

Thompson, F., "Equivalence of Games in Extensive Form," Research Memorandum RM-759, The RAND Corporation, Santa Monica, CA, 1952.

Accepted by Gregory W. Fischer: received August 15, 1993. *This paper has been with the authors* $6\frac{1}{2}$ *months for* 2 *revisions.*