THE KOREAN ECONOMIC REVIEW
Volume 14, Number 2, Winter 1998

DISEQUILIBRIUM BEHAVIOR AND SATISFICING IN THE CENTIPEDE GAME

YOUNGSE KIM*

We reexamine subjects' behavior in the centipede game experiments studied by McKelvey and Palfrey (1992, Econometrica). We assume that players choose acts according to a modified version of case-based decision theory. The calibration results of this paper show that "satisficing" can explain the actual subjects' behavior surprisingly well. More precisely, it is shown that 86 to 92 percentage of the observed behavior is consistent with our model prediction. In our model, the initial aspiration level is not constrained, and the calibration results can be used to evaluate whether the inferred initial aspiration levels are consistent with the payoffs resulting from the subgame perfect equilibrium outcome.

JEL Classification: C72, D83
Keywords: Subgame Perfect Equilibrium, Case-Based Decision, Satisficing, Minimal Inconsistency, Minimal Imperfection

I. INTRODUCTION

Experimental studies make it clear that models of game behavior demanding perfect rationality of the players seem to perform poorly with even very simple games. Consider the situation in which rational players play a finite-horizon extensive form game of perfect information. Traditional theory can generically propose a sharp prediction, called the subgame perfect equilibrium. The trouble is that experimental studies have reported drastically different behavior in some

Received for publication: Feb. 5, 1998. *Revision accepted:* May 18, 1998.

* Department of Economics, College of Economics & Business, Yonsei University, Seoul 120-749, Korea. This research started when I was at London University. I wish to thank Frank H. Hahn and David K. Levine for their advice and encouragement. This paper benefited from discussions with Ken Binmore, Tilman Borgers, Vince Crawford, David Easley, Peter Klibanoff, hamid Sabourian, Dale Stahl and seminar participants at Churchill College Cambridege and the Seventh Econometric Socety World Congress held in Tokyo. Suggestions by an anonymous referee improved the exposition of the paper. Financial support from Yonsei University Research Grant(1996) is gratefully acknowledged. The usual caveats apply.

well-known games, such as ultimatum bargaining games and the centipede game as studied by McKelvey and Palfrey(1992). Any Nash equilibrium to the centipede game results in miserably low individual payoffs despite the possibility of all players becoming much better off. Contrary to this extreme prediction of perfect equilibrium, experimental studies reported that the outcomes in the centipede game were widely dispersed with the mode in the middle.

We propose a model to explain the observed behavior which violates perfect rationality and, in particular, backward induction argument.[1] This paper develops a general theory which can be used to evaluate subjects' behavior and applies the theory to the simple four-move centipede games. The model can be described as follows. Players form their own initial aspiration levels right after being informed about the game but before playing it. This is meant to capture that, without explicitly modeling how, each subject's initial aspiration level is determined by factors such as the informational condition, the first mover advantage or disadvantage, the subject's model of opponents' play, his experience of having faced similar situations before and knowledge about game theory. Once initial aspiration levels are formed, players adapt in subsequent periods to their experience. The dynamic model of this paper is heavily based on Gilboa and Schmeidler's(1995, 1996) case-based decision theory (CBDT henceforth).

In our approach, a player chooses a strategy to maximize the objective function which is the sum of the differences between experienced payoffs and the current aspiration level. The objective function, which Gilboa and Schmeidler derived from a set of axioms, has an element of bounded rationality and habit persistence. The aspiration level evolves over time, which is set to be a weighted average of its previous value and the best average payoff over all strategies. Using Binmore's(1987) terms, in our model the 'eductive' process determines the initial aspiration level, and the 'evolutive' process governs the subsequent learning process. Traditional theory, to verify whether a particular outcome is indeed an equilibrium, often presupposes an introspection process that seems beyond the cognitive limit of human subjects. On the other hand, recent developments in learning and evolution, showing how an equilibrium gets to be played, often assume agents who can do nothing but make occasional adaptations to previous history. Our model may be an approach to incorporate these two features into one framework.

This paper shows that satisficing can explain the subjects' observed behavior surprisingly well. We define the minimally inconsistent path(s) to be the one that, among the paths predicted by the model, is least different from the actual play. Here, 'least' is measured in terms of the number of periods. The

[1] Camerer et al. (1993) used a software system to record the information subjects looked at and to draw inferences about how people think in a ultimatum bargaining game. They concluded that 'any model that assumes players backward induct is a poor descriptive account of how people think. Such a model might make reasonable predictions, but models that incorporate better representations of human cognition could make even better predictions.'(p.45)

calibration results show that the proportion of periods at which the observed behavior is consistent with our model ranges from 86 to 92 percent depending on the payoff size and the player's role. The result does not change significantly regardless of whether the stake is small or large. We also define the minimally imperfect initial aspiration to be the number that, among the initial aspiration levels consistent with any minimally inconsistent paths, is nearest to the perfect equilibrium payoff. If there is some minimally inconsistent aspiration equal to the perfect equilibrium payoff, we can not reject the hypothesis that the subject made his initial choice on the basis of rational expectations. It is shown that roughly half of the subjects' minimally imperfect initial aspirations coincide with the perfect equilibrium payoff. This suggests that our claim would remain true even if half of the subjects can rationally calculate the equilibrium payoff and set their initial aspiration levels accordingly.

It has been recognized that satisficing can capture many salient features of experimental results. Examples include repeated prisoner's dilemma (Selten and Stoecker, 1986), bargaining problems (Mitzkewitz and Nagel, 1993) and a repeated zero-sum game (Mookherjee and Sopher, 1994). The most widely applied framework of satisficing behavior is the stochastic learning mechanism due to Bush and Mosteller(1955) and Harley(1981). In this approach, each player is assumed to adapt his strategy from one round to the next by increasing the probability assigned to the chosen strategy if it resulted in a payoff exceeding the aspiration level, and decreasing it otherwise. Fixing the aspiration level to be zero, this approach presupposes positive payoff in economic term or stimulus in psychological term. Roth and Erev(1995) showed by simulations that the interme- diate term predictions of modified Harley's models track well the observed behavior in three games: a sequential best-shot game, a market game, and an ultimatum bargaining game. Roth and Erev suggest that subjects used essentially the same learning rules regardless of game structure and location and that 'the observed differences reflect different patterns of adaptations'.

We apply a modified version of CBDT on the following ground. Above all, we believe that the environment CBDT presumes is similar to the environment subjects often face in actual laboratory experiments. Gilboa and Schmeidler remark that CBDT is particularly appropriate in analyzing situations involving "ignorance", which refers to the situation in which neither the states of the world nor probabilities on them are naturally defined.[2] In many experimental studies, subjects are asked to make certain decisions, while being placed in unfamiliar situations especially in initial stages. Hence, this paper attempts to test how well CBDT works in situations that we believe are the most pertinent.

[2] Gilboa and Schmeidler added the notion of "ignorance" to the Knightian distinction between risk and uncertainty. To demonstrate how backward induction breaks down in some repeated strategic form games, Dow and Werlang(1994) apply concept of non-additive probabilities, which is a model of uncertainty.

Second, the deterministic nature of our model enables us to check whether and to what extent actually observed paths of subjects' choices are consistent with the theory. This is in sharp contrast to Roth and Erev that relied heavily on simulations, although our approach can convey the same insight and intuition. Last, the informational requirement of CBDT is less demanding relative to other learning models. A case-based player needs to know nothing about others' choices or characteristics, nor the structure of the game at hand. All he must remember are the number of times that each strategy was chosen, if ever tried, and its average performance.[3]

The structure of the paper is organized as follows. The next two sections build up the model and explain central notions for subsequent analysis. Section 4 describes the centipede game and demonstrates how our model works. Section 5 reports the calibration results. The final section concludes.

Ⅱ. THE MODEL

The stage game is a finite two-person extensive form game of perfect information. At each decision node one player decides between a finite number of actions. Let S^i and $S = S^1 \times S^2$ denote the set of player-i's pure strategies and the set of strategy profiles, respectively. Let s^i and s denote typical elements of the sets S^i and S, respectively. Each strategy profile determines a terminal node $z \in Z$ giving a vector of payoffs, one to each player. Let $\pi^i : Z \rightarrow \Re$ denote player i's payoff function. Also let $\underline{\pi}^i$ and $\overline{\pi}^i$ denote the lowest and highest payoff achievable in the stage game for player i, respectively. It is well known that this class of games has generically a unique subgame perfect equilibrium. Let Π^i_{spe} denote the set of payoffs that player i receives on the subgame perfect equilibrium.

According to game theory tradition, a strategy is defined as a complete contingent plan. In belief-based learning approaches, it is important to specify off-the-equilibrium beliefs and plays. However, in our model of bounded rationality, it is more sensible to assume that players do not distinguish between two pure strategies that lead to the same terminal node. For this reason, a "strategy" in this paper refers to a pure strategy that remains after realization equivalent strategies are merged in an equivalent class.

Imagine that the stage game is repeatedly played at periods $t = 1, 2, \cdots, T$, where T may be finite or infinite. Player i updates his strategy at period t, according to player i's payoff realized up to $(t-1)$, relative to his present

[3] Belief formation models, such as Fudenberg and Kreps(1992) and Canning(1992), require that players observe the empirical frequencies of opponents' realized strategies in order to form beliefs against which to choose a best response. Evolutionary dynamics, such as Kandori, Mailath and Rob(1993) and Young(1993), require that the aggregate characteristic of the population be observable.

aspiration level. Let H_t^i denote player i's aspiration level at the beginning of the period t. The set of T-histories is a subset of $\Omega' \equiv (\mathfrak{R}^2 \times S \times \mathfrak{R}^2)^T$. A T-history $\omega = \{H_t, s_t, \pi_t)|1 \leq t \leq T\} \in \Omega'$ will be interpreted as follows: for all $t \geq 1$, the pair of aspiration levels is $H_t = (H_t^1, H_t^2)$ at the beginning of the period, a strategy profile $s_t = (s_t^1, s_t^2)$ is chosen, and it yields a payoff vector of $\pi_t = (\pi^1(s_t), \pi^2(s_t))$. The projection functions $\sigma_t^i : \Omega' \to S^i$ and $\sigma^i : = (\sigma_t^i)_{1 \leq t \leq T}$ have the obvious meaning.

Next we define a function $C^i : \Omega' \times S^i \times T \to 2^T$ to be the set of periods, up to a given time, at which player i played a given strategy, according to a given history. That is,

$$C^i(\omega, s^i, t) = \{\tau < t | s_\tau^i(\omega) = s^i\}.$$

We also define the number of times player i played a particular strategy s^i by

$$K^i(\omega, s^i, t) = \# \, C^i(\omega, s^i, t) \in \{0, 1, 2, \cdots, T\}.$$

The strategy is evaluated by the sum of the differences between experienced payoffs and the current aspiration level. Formally, player i's objective at period t is to choose a strategy that maximizes the following U^i-function:

$$U^i(\omega, s^i, t) = \begin{cases} 0, & \text{if } K^i(\omega, s^i, t) = 0 \\ \sum_{\tau \in C^i(\omega, s^i, t)} [\pi_\tau^i(\omega) - H_t^i(\omega)], & \text{if } K^i(\omega, s^i, t) > 0 \end{cases} \tag{1}$$

where $\pi_\tau^i(\omega)$ is the player i's payoff at period t along the history ω. If the U^i-maximizing strategies are multiple, then player i is assumed to choose one of the strategies by some deterministic tie-breaking rule. Notice that the value U^i directly depends only on the current aspiration level H_t^i. It will be often convenient to express the U^i function in the following form:

$$U^i(\omega, s^i, t) = \begin{cases} 0, & \text{if } K^i(\omega, s^i, t) = 0 \\ K^i(\omega, s^i, t)[\overline{\Pi}^i(\omega, s^i, t) - H_t^i(\omega)], & \text{if } K^i(\omega, s^i, t) > 0 \end{cases} \tag{2}$$

where $\Pi^i(\omega, s^i, t) = \dfrac{\sum_{\tau \in C^i(\omega, s^i, t)} \pi_\tau^i(\omega)}{K^i(\omega, s^i, t)}$ is the absolute average payoff of the strategy s^i.

Now consider the aspiration revision rule. For a given adjustment weight $\alpha^i \in [0, 1]$ and the initial aspiration level $H_1^i \equiv H^i \in [\underline{\pi}^i, \overline{\pi}^i]$, player i updates his aspiration level in an adaptive manner. Specifically, the present aspiration level is the weighted average of its previous value and the maximal average payoff, where the maximum is taken over all strategies. On the other hand, if player i had no opportunity to move at $(t-1)$, then the resulting stage game payoff is used in revising the next period aspiration level. More formally,

$$H_t^i = (1-a^i)H_{t-1}^i + a^i \max_{s^i \in S^i(t)} \{\overline{\Pi}^i(s^i, t)\}, \tag{3}$$

where $S^i(t)$ is the set of player i's strategies which were chosen before t. This revision rule has the desirable property that if the same strategy profile is repeatedly played then the aspiration level approaches the realized payoff.

Ⅲ. THE ANALYSIS

In this section we explain how the theoretical framework just developed can be used to evaluate observed behavior. Experimental data at individual level provide the sequence of realized outcomes. Since there is a set of strategies which is consistent with a given outcome, corresponding to a sequence of realized outcomes there is the set of the sequences of strategies that is consistent with the actual data. We call it as the set of actual paths. Let $\hat{S}^i$ and $\hat{S}^i = (\hat{S}_t^i)_{1 \le t \le T}$ be the set of actual paths for player i and its typical element, respectively.

Now we want to characterize the paths of strategies that are consistent with our model, namely the sequences of strategies that a player who behaved as if he was a modified CBDT decision-maker would have chosen. To this end, fix the initial actual choice $\hat{S}_1^i$ as given. We can then generate the subspace of T-histories that, starting from the initial strategy $\hat{S}_1^i$, are compatible with U^i -maximization and aspiration revision rule. Formally,

$$\begin{aligned} \Omega = \{\omega \in \Omega' | s_1^i(\omega) = \hat{s}_1^i,\ s_1^i(\omega) \in \textbf{\textit{argmax}}\ U^i(s^i, t),\ \forall t \ge 2, \text{ and} \\ \exists\ a^i \in (0,1) \text{ such that } (H_t^i)_{1 \le t \le T} \text{ evolves by Eq.(3)}\} \end{aligned} \tag{4}$$

Focus on a particular subject playing player-i's role, so we suppress the superscript i whenever there is no confusion.

We say that a play of player i is *consistent with the model* if $\hat{s}^i \in \{\sigma^i(\omega) | \omega \in \Omega\}$, where $\sigma^i(\omega)$ is the projection function. Suppose that the observed play of a particular player is not perfectly consistent with the model. We postulate that his intended play is the path minimizing the number of periods at which he behaved inconsistently. More formally,

Definition 1. Define the set $\Omega^*(\hat{s}^i)$ to be the set of *minimally inconsistent* (MIC, in short) histories that is associated with the actual path $\hat{s}^i$, where

$$\Omega^*(\hat{s}^i) \equiv \{\omega \in \Omega \mid \#\{t \mid \hat{s}^i_t \neq \sigma^i_t(\omega)\} \leq \#\{t \mid \hat{s}^i_t \neq \sigma^i_t(\omega')\},\ \forall \omega' \in \Omega\}. \tag{5}$$

Let $\Omega^* \equiv \bigcup_{\hat{s} \in \hat{S}} \Omega^*(\hat{S})$ be the set of all MIC histories. Define $\sigma^i[\Omega^*(\hat{s}^i)]$ to be the set of MIC paths that is associated with the actual path $\hat{s}^i$, where $\sigma^i(.)$ is the projection function from histories to strategy choices. Clearly, $\Sigma^i := \bigcup_{\hat{s} \in \hat{S}} \sigma^i[\Omega^*(\hat{s}]$ is the set of all MIC paths.

Corresponding to each MIC path, there exist pairs of the initial aspiration level and aspiration revision coefficient, (H, α), that are consistent with the given path. Let a minimally inconsistent aspiration level (MICA, in short) be the initial aspiration level which is compatible with some MIC path for some revision coefficient, α.

The following story is what we keep in mind about how people behave. Subjects form their initial aspiration levels after being informed of the game but before ever playing it. We do not attempt to model how players form their initial aspirations. We believe that un-modeled factors, such as rationality of players, the first mover (dis)advantage, whether the player has faced a similar situation before, affect the level of initial aspiration. For this reason, if there is some MICA equal to the perfect equilibrium payoff, we may not be able to reject the hypothesis that the subject made his initial choice on the basis of rational expectations. Hence, we are interested in the MICA that is nearest to the perfect equilibrium payoff. We call it to be the minimally imperfect initial aspiration level (MIPA, in short) and the MIC paths associated with MIPA to be the minimally imperfect (MIP, in short) paths.

We want to formally define these notions. Let $\Lambda^1 = \{H^1(\omega) \in [0.4, 6.4] \backslash \omega \in \Omega^*\}$ and $\Lambda^2 = \{H^2(\omega) \in [0.1, 3.2] \backslash \omega \in \Omega^*\}$ be the set of player 1's and 2's initial aspiration levels that are compatible with some MIC path. Now the notion of minimal imperfection is defined as follows.

Definition 2. Define the *minimally imperfect initial aspiration level*(MIPA, in short) to be:

$$\lambda^i = \arg \inf_{\substack{\pi^i \in \Pi^i_{spe} \\ H^i \in \Lambda^i}} |\pi^i - H^i|, \tag{6}$$

where Π^i_{spe} is the set of perfect equilibrium payoffs to player-i and the argument is taken over H^i's. Also define a *minimally imperfect path* to be an element of the set $\{s^i \in \Sigma^i \mid H^i(\omega) = \lambda^i\}$ where $H^i(.)$ is the projection function from histories to initial aspiration levels.

We will later demonstrate how to check consistency and calculate the minimal imperfect initial aspiration level and the corresponding path, using an individual datum in the centipede game.

Ⅳ. THE CENTIPEDE GAME

The Four-move Centipede Game. Consider the centipede game studied by McKelvey and Palfrey (1992, MP henceforth). Figure 1 illustrates the four-move extensive form. An experimental treatment was conducted with stakes four times those indicated above. This experiment is labeled "high payoff". Each player has three strategies in our sense, namely T_i, P_iT_{i+2} and P_iP_{i+2}, for $i=1,2$. As was previously noted, although player-i has four pure strategies, T_iT_{i+2}, P_iT_{i+2}, T_iP_{i+2} and $P_iP_{i=2}$, for $i+1,2$, we consider two strategie T_iT_{i+2} and T_iP_{i+2} as one strategy.

Since any Nash equilibrium to this game involves player-1 playing T_1, the unique subgame perfect equilibrium clearly makes the same prediction. That is, $\Pi^1_{spe}=\{0.4\}$ and $\Pi^2_{spe}=\{0.1\}$ Experimental outcomes are quite different from the Nash prediction. Each player played the game nine to ten times against different opponents. MP's aggregate statistics found only 37 out of 662 games ending with T_1, while 23 of the games end with both players passing at every move. The remainder of the outcomes are scattered in between.

A Heuristic Example. We demonstrate how to calculate the set of MIC paths and the MIPA. The datum is for the third subject of player 1's role in MP experimental treatment Session 2.[4] Throughout the example, we suppress the index i and the dependence of aspiration level on the adjustment coefficient, α, and the initial aspiration level, H. The first, second and fourth row of Table 1 respectively indicate the index of round, the observed terminal node in each round and the realized payoff to the particular subject at hand. From the series

[Figure 1] The Four-Move Centipede Game

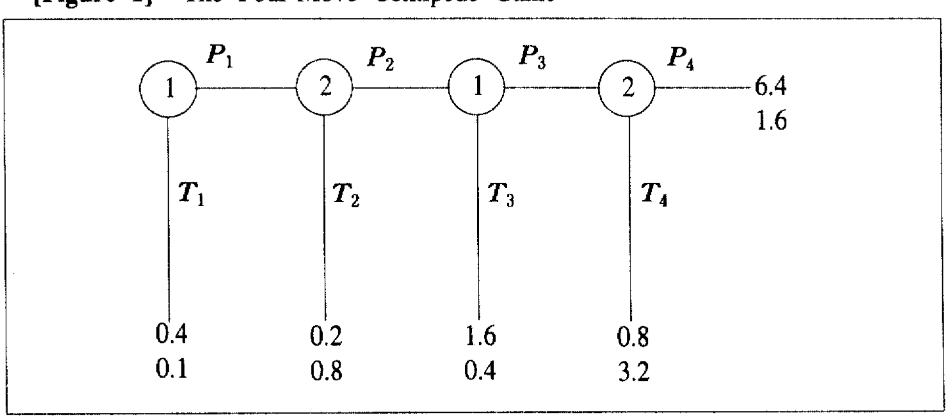

[4] The calculations are purely mechanical. We do not suggest at all that the chosen subject's behavior is particularly interesting.

[Table 1] A Heuristic Example π_t

t	1	2	3	4	5	6	7	8	9
Z	3	3	2	2	1	1	1	1	1
$\hat{s}^i$	PT	PT	P •	P •	T	T	T	T	T
π_t	1.6	1.6	0,2	0.2	0.4	0.4	0.4	0.4	0.4
path A	PT	PT	PT	PT	T	T	T	T	T
map.	1.6	1.6	3.4/3	0.9	0.9	0.9	0.9	0.9	0.9
path B	PT	PT	PT	T	T	T	T	T	T
map	1.6	1.6	3.4/3	3.4/3	3.4/3	3.4/3	3.4/3	3.4/3	3.4/3
path C	PT	PT	PP	T	T	T	T	T	T
map	1.6	1.6	1.6	1.6	1.6	1.6	1.6	1.6	1.6
path D	PT	PT	PP	T	T	T	T	T	T
map	1.6	1.6	1.6	1.6	1.6	1.6	1.6	1.6	1.6

Note: The "map" means the maximal average payoff.

of realized terminal nodes we can figure out the corresponding series of strategies which are consistent with the observed play. This is shown on the third row. Notice that there are four possible paths depending on $\hat{s}_3 \in \{PT, PP\}$ and $\hat{s}_4 \in \{PT, PP\}$.

Consider the first possible path, labeled as 'path A', in which $(\hat{s}_3, \hat{s}_4) = (PT, PT)$. The sixth row, labeled as 'map A', of Table 1 shows the sequence of maximal average payoff along this path. In order for PT to be consistent from $t=2$ to $t=4$, it must be that

$$PT = \arg\max_{\{T, PT, PP\}} \{0, 1.6 - H_2, 0\},$$

$$PT = \arg\max_{\{T, PT, PP\}} \{0, 3.2 - 2H_3, 0\},$$

and

$$PT = \arg\max_{\{T, PT, PP\}} \{0, 3.4 - 3H_4, 0\},$$

For the switch from PT to T at $t=5$ to be consistent with the model, it must be that

$$T = \arg\max \{0, 3.6 - 4H_5, 0\},$$

It is straightforward to check that plays from $t=6$ to $t=9$ are consistent, regardless of the parameter values. Hence, this subject's play is perfectly consistent with the model if there exists (a, H) satisfying the following two conditions:

$$H_4(\alpha,\ H) \leq \frac{3.4}{3} \text{ and } H_5(\alpha,\ H) \geq 0.9 \tag{7}$$

It is easy to show that there exists such $(\alpha,\ H)$ satisfying the above inequalities. This path is perfectly consistent with the theory, and is thus minimally inconsistent. Among those Hs satisfying Eq.(7), the value $H=0.9$ is the nearest to the equilibrium payoff, 0.4. Therefore, the MIPA is 0.9.

Consider the second possible situation, labeled as 'path B', in which The corresponding sequence of maximal average performances along path B appears on the eighth row, labeled as 'map B', of Table 1. In order for $\hat{s}_2 = PT$ to be consistent, it must be that

$$PT = \underset{s_2 \in \{T,\ PT,\ PP\}}{\arg\max} \ \{0,\ 1.6 - H_2,\ 0\}.$$

Likewise,

$$PT = \underset{s_3 \in \{T,\ PT,\ PP\}}{\arg\max} \ \{0,\ 3.2 - 2H_2,\ 0\},$$

$$PP = \underset{s_4}{\arg\max} \ \{0,\ 3.4 - 3H_4,\ 0\},$$

$$T = \underset{s_5}{\arg\max} \ \{0,\ 3.4 - 3H_5,\ 0.2 - H_5\},$$

and

$$T = \underset{s_t}{\arg\max} \ \{(t-5)(0.4 - H_t),\ 3.4 - 3H_t,\ 0.2 - H_t\}, \text{ for } t = 6, 7, 8, 9$$

It is easy to show that there cannot exist a pair of parameters, $(\alpha,\ H) \in [(0,1) \times [0.2,\ 6.4]]$ such that strategy T maximizes the objective function at $t=7,\ 8$ and 9. We can also show that all other five equations corresponding to the period $t=2$ through $t=6$ hold for $(\alpha,\ H)$ pairs in the area encompassed by three equations:

$$\alpha = 0, \tag{8}$$

$$H = 1.6, \text{ and} \tag{9}$$

$$H_6(\alpha,\ H) \equiv H(1-\alpha)^5 + 1.6\alpha(1-\alpha)^3[(1-\alpha)+1] + \frac{3.4}{3}\alpha[(1-\alpha)^2 + (1-\alpha) + 1] = 1.5. \tag{10}$$

Since the smallest value of H satisfying the above three equations is 1.5, we conclude that the MIPA is 1.5.

There remain two possibilities of actual path, labeled C and D in Table 1. It is tedious but straightforward to show that, for each of those two paths, the number of inconsistent plays is non-null. Therefore, we conclude that this subject's play is perfectly consistent with the theory along the singleton set of MIC path, $\hat{s}=\{PT, PT, PT, PT, T, T, T, T, T\}$, and that the MIPA is 0.9.

V. CALIBRATION RESULTS

Table 2 reports the calibration results.[5] The first column indicates the number of periods at which the observed path is distinct from some MIC paths. Reading downwards, each cell reports the number of subjects whose play was inconsistent for the corresponding number of periods. Consider the low payoff case. Our model perfectly predicts plays of 28 subjects out of 58, and 51 subjects except at most two periods. The last row of the table indicates the aggregate proportion of periods at which the observed plays are different from MIC paths. It is shown that 88.3 percentage of player-1s' plays and 89.7 percentage of player-2s' plays are consistent with the model.

The results with high payoffs are not significantly different. At the aggregate level, 86 percentage of player-1s' plays and 92 percentage of player-2s' plays are consistent with the model. More precisely, we test the goodness-of-fit of two frequency distributions. The null hypothesis is that the frequency distribution with high payoffs fits well the low payoff counterpart. The chi-square statistics are 2.01 for the player-1s and 1.22 for the player-2s, both of which are lower than the critical value 7.81 at the 5% significance level. This implies that there is good agreement with the null hypothesis.

[Table 2] The Results of Minmal Inconsistency

	Player ls		**Player 2s**	
	Low payoff treatment	**High payoff treatment**	**Low payoff treatment**	**High payoff treatment**
None	**14**	**6**	**14**	**5**
1	**5**	**0**	**5**	**2**
2	**6**	**2**	**7**	**3**
3 or more	**4**	**2**	**3**	**0**
%	**11.7**	**14.0**	**10.3**	**8.0**

Notes. 1. The total number of subjects is 78, among whom 58 are for the low-payoff treatment and the remaining 20 are for the high-payoff treatment.

2. The first column indicates the number of periods at which the observed path is minimally inconsistent with the model. Each cell indicates the number of subjects whose play is consistent EXCEPT the corresponding number of periods.

3. The last row indicates the aggregate proportion of periods in which the actual path different from the minimally inconsistent paths.

[5] Computer programs and worksheets are available upon request.

Table 3 summarizes the frequency of MIPAs. The second row of the table indicates the number of player 1s (respectively, player 2s) whose MIPA is exactly equal the perfect equilibrium payoff, 0.4 (resp., 0.1). The third column indicates the number of subjects whose MIPAs lie between 0.2 and 0.8 (resp., 0.1 and 0.4). The last column indicates the number of subjects whose MIPAs are no larger than 1.6 (resp., 0.8), which corresponds to one-quarter of the total stake. Twenty two subjects out of 39 player-1s have their MIPAs equal to the perfect equilibrium payoff ($0.4), while fifteen out of 39 player-2s have their MIPAs equal to the perfect equilibrium payoff ($0.1). All except one of player-1's have their MIPAs no larger than the quarter of the total stake, and 30 out of 39 subjects of player-2's role have their MIPAs no larger than the quarter of the stake.

If a particular subject's MIPA is equal to the perfect equilibrium payoff, we can not reject the hypothesis that the subject made his initial choice on the basis of rational expectations. The results on MIPA show that roughly half of the subjects' MIPAs coincide with the perfect equilibrium payoff. This suggests that half of the subjects might rationally calculate the equilibrium payoff and set their initial aspiration levels accordingly. By definition, the MIPA is the aspiration level that is made as tight as possible to the equilibrium payoff. This means that the actual initial aspiration levels are widely dispersed, which plays a crucial role in the behavior not converging to the equilibrium prediction.

[Table 3] Summary Frequency of MIPAs $\pi_{spe}=0.1$

Player-1s		Player-2s	
Equal to $\pi_{spe}=0.4$	22	**Equal to $\pi_{spe}=0.1$**	15
Between 0.2 and 0.8	27	**Between 0.1 and 0.4**	23
Between 0.2 and 1.6	38	**Between 0.1 and 0.8**	30
Total number of subjects	39	**Total number of subjects**	39
Minimum payoff	0.2	**Minimum payoff**	0.1
Maximum payoff	6.4	**Minimum payoff**	3.2

Note. 1. All low and high treatments are aggregated.

An informal description of intuition follows. Assume that player 1s' initial aspiration levels are widely distributed over the interval [0.2,6.4], while those of player 2s' are distributed over [0.1,3.2]. Around the top of the game tree, the payoff from TAKE are so small that most subjects do not satisfy. Player 1's choosing T_1 would yield the payoff 0.4, which is very close to the lower bound of the interval [0.2,6.4]. Similarly, player 2's choosing T_2, given the chance to move, would yield 0.8 which is around the quartile of the initial support [0.1,3.2]. Thus, the satisficing players would have extremely negative bias against the strategies prescribed by the subgame perfect equilibrium. Around the root of the game tree, on the other hand, backward induction stops exerting a considerable influence, since players stay with a sub-optimal strategy as long as it fares better than their current aspirations. Both effects work together so as to make the outcomes disperse over all terminal nodes. As suggested above, this remains true even if half of the subjects can figure out the game correctly and set their initial aspiration levels accordingly.

The salient feature of time-series data is that outcome distributions tend to be more inclined to the subgame perfect outcome as subjects gain more experience. Frequencies at the first terminal node Z_1 and the last terminal node Z_5 both fluctuate between 0 and 0.103. Frequencies at node Z_2 increase from 0.276 to 0.483 as time passes, whereas those at node Z_4 decrease from 0.241 to 0.050. The terminal node Z_3 in the middle shows a widely fluctuating frequency from 0.310 to 0.448 over time. We identify this trend as a gradual movement towards subgame perfection, though not necessarily a convergence. Player 1s are more likely to be frustrated by choosing the strategy P_1P_3 as time passes since player 2s eventually learn to choose T_4 at the final decision node unless his aspiration level is very low. Some player 1s learn that the strategy P_1P_3 did not perform satisfactorily in previous matches. Note that this scenario does not necessarily mean that players eventually learn to backward-induct, since there can be those player 2s who could not confirm the fact that strategy P_2T_4 is better than P_2P_4. MP observed that there are several subjects who pass at every opportunity they have. In our framework, a player may always pass if his initial aspiration level is quite low and has not confirmed the fact that a better strategy than PP is available.

VI. CONCLUDING REMARKS

Our model is deterministic. Many seminar participants and referees suggest that the right direction is to allow tie-breaking randomization and to estimate the maximum likelihood parameters. Although I agree this MLE method makes another sense, I intentionally did not follow their advice. The reason is that, from the traditional theorist's viewpoint, the decision-maker in our model may be hard to accept. Decision-makers maximize an objective function that incorporates aspiration and cumulative payoffs, but not the average realized payoff. Aspiration evolves in an adaptive manner but not in a forward-looking way. If we allow players to randomize or tremble, the parsimony of the model might be extremely problematic.

Besides the aspiration update, we could have introduced another type of dynamic, namely expansion of the set of choices. Roughly speaking, if strategies that were available resulted in unsatisfactory outcomes, the decision-maker begins searching for new available strategies, which leads to an enlarged set of choices.

If we modify the model in this way, it turns out that the set of available strategies, S^i, can be treated as a choice variable of the outside analyzer who is free to choose S^i to make the data fit the theory.7 To put it polemically, if some player would always choose the same strictly dominated strategy this could be justified as being completely rational under the approach pursued here by simply assuming that the player fails to realize that other strategies are available. However, in order to let the model parsimonious, we do not formalize this idea.

One obvious shortcoming of this paper directly stems from lack of data. The existing experiments were conducted for nine to ten periods, on which our analysis of consistency are based. This is clearly not sufficient for a study of dynamic processes. The high performance of our model might be even due to the fact that it fits only ten-period time series data.

A number of questions are raised about works that are currently being undertaken, and problems for future research. There are other competing theories available. A model of bounded rationality is the Bush-Mosteller(1955) type stochastic learning approach. There are also attempts to reconcile experimental data with more traditional game-theoretic predictions. Camerer and Weigelt(1988) and McKelvey and Palfrey(1992) use the home-made priors approach, but such a method seems to difficult to generalize to other games. Fudenberg and Levine (1997) argue that, using the self-confirming equilibrium notion introduced by Fudenberg and Levine(1993), the proportion of players that would need to have irrational payoffs to generate the observed path is small. They report that the average loss of a player is $0.03 to $0.64 in a game involving stakes between $2 and $30. From a methodological point of view, while Fudenberg and Levine's analysis are based on the aggregate probability distribution over outcomes, we consider the complete history of each individual player's choices.

It will be interesting to compare which approach fares best in applicable experimental data. The companion paper(1995) exhibits high performance of our model in the ultimatum bargaining game and the sequential best-shot game. Again, a critical limitation is that the existing experimental studies provide very short periods data. An agenda for future research is to conduct laboratory experiments for a much longer periods and to test various theories by using new experimental data.

REFERENCES

Binmore, K. G. (1987), Modeling rational players, I and II, *Economics and Philosophy*, 3, 179-214, and 4, 9-55.

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

Camerer, C., Johnson, E. J., Rymon, T., and Sen, S. (1993), Cognition and framing in sequential bargaining for gains and losses, In Binmore, K. G., Kirman, A., and Tani, P. (eds.), *Frontiers of Game Theory*. MIT Press, pp.27-47.

Camerer, C., and Weigelt, K. (1988), Experimental tests of a sequential equilibrium reputation models, *Econometrica*, 56, 1-36.

Canning, D. (1992), Average behavior in learning models, *Journal of Economic Theory*, 57, 442-72.

Dow, J., and Werlangm, S. R. C. (1994), Nash equilibrium under Knightian uncertainty: breaking down backward induction, *Journal of Economic Theory*, 64, 305-24.

Fudenberg, D., and Kreps, D. (1992), *Lecture on Learning and Equilibrium in Strategic Form Games*, CORE Lecture Series.

Fudenberg, D., and Levine, D. K. (1993), Self-confirming equilibrium, *Econometrica*, 61, 523-45.

__________ (1997), Measuring players' losses in experimental games, *Quarterly Journal of Economics*, 112, 507-36.

Gilboa, I., and Schmeidler, D. (1995), Case-based decision theory, *Quarterly Journal of Economics*, 605-39.

Gilboa, I., and Schmeidler, D. (1996), Case-based optimization. *Games and Economic Behavior*, 15, 1-26.

Harley, C. B. (1981), Learning the evolutionarily stable strategies, *Journal of Theoretical Biology*, 89, 611-33.

Kandori, M., Mailath, G., and Rob, R. (1993), Learning, mutation and long-run equilibria in games, *Econometrica*, 61, 27-56.

Selten, R., and Stoecker, R. (1986), End behavior in sequences of finite prisoner's dilemma supergames: a learning theoretic approach, *Journal of Economic Behavior and Organization*, 7, 47-70.

Simon, H. A. (1955), A behavioral model of rational choice, *Quarterly Journal of Economics*, 69, 99-119.

Young, P.H. (1993), The Evolution of conventions, *Econometrica*, 61, 57-84.