GAMES AND ECONOMIC BEHAVIOR **10**, 218–254 (1995)

# On Players' Models of Other Players: Theory and Experimental Evidence*

DALE O. STAHL AND PAUL W. WILSON

*Department of Economics, University of Texas, Austin, Texas 78712*

Received May 25, 1994

We develop and test a theory of human behavior in 3 × 3 symmetric games. The theory hypothesizes a family of five boundedly rational archetypes distinguished by their model of other players and their ability to identify optimal choices given their priors. We designed and conducted an experiment to detect these archetypes as well as a rational expectations type and to estimate parameters which define these types. The experimental evidence rejects the rational expectations type but confirms the boundedly rational theory. We consider this a stepping stone toward a descriptive and prescriptive theory of games. *Journal of Economic Literature* Classification Numbers: B41, C14, C51, C72, C90. © 1995 Academic Press, Inc.

## 1. INTRODUCTION

It has been known for some time that game theory is a poor predictor of human behavior in experimental settings.[1] A major step toward a descriptive theory was made by McKelvey and Palfrey (1992) who added a specific error component to the pure theory and conducted maximum likelihood estimation of the error parameters. Their model posits two

* Partial funding of this research was provided by Grants SRG-233 and SRG-427 from the University of Texas Research Institute and by Grant SBR-9308914 from the National Science Foundation. We thank Xiaohua Lu for research assistance and Ray Battalio, John van Huyck, Tom Palfrey, Reinhard Selten, the participants of the 1993 Workshop on Experimental Game Theory at Stony Brook, and the participants of the 1994 Conference on Learning Equilibrium Behavior at Texas A&M University for encouragement. We are also grateful for comments from anonymous referees. Of course, any remaining errors are our responsibility.

[1] For example, see Camerer and Weigelt (1988), Ochs and Roth (1989), Cooper *et al.* (1990), Brandts and Holt (1992), and Van Huyck *et al.* (1990, 1991, 1992).



0899-8256/95 $12.00

types of players: altruists and Bayesian rational players. Nagel (1993) proposed and presented evidence for a hierarchical model of bounded rationality. In Stahl and Wilson (1994), we presented a model with several boundedly rational types (based on Nagel's model) and some supporting experimental evidence.[2] The theory we present here is considerably evolved, the experimental design is improved, and the statistical testing is more extensive. Our current multi-modal theory includes a hierarchical conception of bounded rationality that is captured in a parsimonious econometric specification, and it also allows for players with fully rational expectations. Moreover, we test the out-of-sample predictability of our model.

A crucial concept in the theories of rational choice under uncertainty (Savage, 1954; and Anscombe and Aumann, 1963) is the decision-maker's subjective probability assessment of the possible events not under his direct control. These probability assessments come from *models* of these events. For instance, in the case of a die, it is customary to assume a 1/6-th chance of any face, implicitly invoking the standard model of a throw of a fair die.

To extend this theory to multi-decision-maker problems (i.e., games), we need "only" specify for each player his *model of the other players*. However, this is a difficult task, especially when we insist that all players are perfectly rational, have identical models of other players, and that these models are observationally equivalent to the true model (whatever that may be). Our task poses an infinitely regressive self-referential problem: a model of a rational player must include a model of rational players which must include a model of . . . . Mathematically, we seek a fixed point in some "model space."[3] Binmore (1987) shows that this self-referential problem is unsolvable within the class of models that could be represented as universal Turing machines. For those readers who do not have faith that human beings possess powers beyond any universal Turing machine, we offer an alternative theory. Our approach builds on Binmore's suggestion that players might truncate an internal simulation of the self-referential problem, and it also grows out of Stahl's (1993) hierarchical model of intelligence.

Decision theory provides us with our starting point. Players need to (1)

[2] In a recent paper, El-Gamal and Grether (1993) test a multi-modal theory for a *single decision-maker* task and identify distinct types of behavior. Holt (1993) finds evidence for different types of learning behavior in experimental data on coordination games.

[3] Viewing each player's subjective belief about how the other players will behave as shorthand for such a model, we can interpret Nash equilibrium as a solution (at least for the class of finite games with a unique Nash equilibrium). On the other hand, when there are multiple Nash equilibria, the "model" is incompletely specified and the problem remains unsolved.

form priors about the behavior of other players and (2) choose best responses given these priors. Using this framework, there are two basic ways players can differ: (1) in their prior (or their model of other players),[4] and (2) their abilities to identify best responses.

The ability to identify best responses can be represented by a standard discrete choice model, in which the payoff from each strategy is computed with some error. Assuming an appropriate error distribution, one obtains a logit probability function for the choices conditional on the prior. The precision parameter of the error identifies one member of the family of probabilistic choice functions. One archetype would be a player with zero precision, in which case the probabilistic choice function would be uniform; we call these *level-0 types*.

The prior of a player entails the player's model of other players. If a player believes that all other players are level-0 types, then his prior will be uniform; we call such players *level-1 types*. If a player believes that all other players are level-0 and level-1 types, then his prior will be a convex combination of the uniform distribution and an average level-1 logit choice function; we call such players *level-2 types*.

We could continue this hierarchy of types indefinitely; however, we recognize that the benefits of higher levels of this hierarchy can decrease substantially. For example, if half of the population is level-2 or below, then the choice of the next level of the hierarchy will rarely differ from the level-2 choice. Further, if higher levels incur higher maintenance costs, then evolution can be consistent with most of the population consisting of fairly low-level types (Stahl, 1993). Given these reasons and the fact that higher levels are unlikely to be identifiable in $3 \times 3$ games, we truncate the hierarchy at level-2.

For the purposes of hypothesis testing of alternative theories, it is necessary to construct an encompassing econometric model. In particular, we want to allow for the behavior suggested by game theory, namely, Nash equilibrium behavior. Considering only games with a unique symmetric Nash equilibrium and interpreting it as a belief (i.e., prior) about other players, a *naive Nash-type* player would have the Nash equilibrium as his prior and compute a best response (possibly with error).

Given that some players may understand the Nash equilibrium concept while others may not, it is reasonable to propose a type that believes some of the players are naive Nash types while others are level-0, level-1, and level-2 types. We call these types *worldly*.

[4] That individuals may differ in their priors is also a major point of Roth and Malouf (1979).

An unboundedly rational player (the subject of traditional game theory) when facing a population of these boundedly rational types would have a prior that is equal to the expected choice probabilities of these types plus the anticipated choice probabilities of other unboundedly rational players. Call such players *rational expectations (RE) types*.

In Section 2 we develop the formal theory of the boundedly rational types complete with parameter specifications. Each type is identified by a parameter triplet such that, given any symmetric game, the model generates the predicted probability distribution of choices. We also develop the formalities of the RE types conditional on the population distribution of the three-parameter family of boundedly rational types and RE types. A well-specified econometric model is a rare opportunity for experimental game theory.[5]

Since we want to estimate these parameters from experimental data, the experimental design entails a variety of symmetric games. We want to focus on the initial choices of our players rather than on "learned" behavior because these initial responses are more difficult to predict, exhibit more diversity, are crucial to whatever learning follows, and can be potentially revealing about the "model of the world" that players bring to a new strategic setting.[6] Hence, the design has each player playing each game exactly once with no feedback between games. More experimental details are given in Section 3.

Our experimental data, presented in Section 4, allows us to estimate the parameters of the boundedly rational types and RE types and to test for their presence in the population. Section 5 presents the statistical analysis. The analysis suggests that RE types are not present in our sample, but confirms the theory for the boundedly rational types. We also compute semi-parametric Bayesian posteriors for the type of each participant and find that most posteriors put 90% or more probability on one type, suggesting that the participants' behavior arises from one "model of other players" for all the games.

In Section 6 we present the results of several robustness exercises based on estimating our model on a subset of the games and predicting the behavior on the remaining games. Both the behavioral predictions and the posterior type identifications are remarkably robust. Section 7 discusses our findings.

[5] Notable exceptions are McKelvey and Palfrey (1992), El-Gamal and Palfrey (1995), and El-Gamal *et al.* (1993, 1994).

[6] Learning and long-run behavior is certainly very important, but so is the initial behavior, and there has been relatively little attention devoted to it. After all, game theory makes predictions about one-shot games, and the "experience" of the players is held to be irrelevant.

## 2. Development of the Formal Theory

### 2.1. *The Three-Parameter Family of Boundedly Rational Types*

Consider a symmetric two-player game with three strategies. Let $j \in \{1, 2, 3\}$ denote a strategy, and let $U_i$ denote the $3 \times 3$ matrix of expected utility payoffs to a row player for game $i \in \{1, 2, \ldots, 12\}$; $U_{ijk}$ is the payoff to a row player in game $i$ when that player chooses strategy $j$ and the column player chooses strategy $k$.

Let $p \in \Delta \equiv \{x \in \Re^3_+ | \Sigma_j x_j = 1\}$. Then, for game $i$, $U_i p_i$ is the $3 \times 1$ vector of expected utility payoffs and $U_{ij} p_i$ is the expected payoff of strategy $j$ facing a population whose distribution across strategies is $p_i$ in game $i$.

For notational convenience, let $P_0$ denote the uniform distribution over $\{1, 2, 3\}$. We then let $b(P_0)$ denote the best response to the uniform distribution and $b^2(P_0) \equiv b[b(P_0)]$ denote the best response to the best response to the uniform distribution. Also, let $p_i^{\text{NE}} \in \Delta$ denote the Nash equilibrium (NE) strategy vector for game $i$.

Define

$$q_{ij}(\mu, \varepsilon) \equiv \varepsilon \frac{\exp(\mu U_{ij} P_0)}{\Sigma_k \exp(\mu U_{ik} P_0)} + (1 - \varepsilon) p_{ij}^{\text{NE}}, \tag{1}$$

and

$$y_{ij}(\mu, \varepsilon) \equiv U_{ij} q_i(\mu, \varepsilon). \tag{2}$$

Equation (1) specifies a two-parameter family of priors. If $\varepsilon = 0$, then we have the Nash prior. If $\varepsilon = 1$ and $\mu = 0$, then we have the level-1 (i.e., uniform) prior. If $\varepsilon = 1$ and $\mu > 0$, then we have the level-2 prior (an imprecise best response to $P_0$). Note that for the level-2 prior, we could also have a term for the level-0 types, but we have found that we cannot empirically identify that parameter because a convex combination of the uniform distribution and a level-1 distribution is essentially the same as a level-1 distribution with a lower precision. Equation (2) gives the expected payoff, $y_{ij}(\mu, \varepsilon)$, to strategy $j$ in game $i$ conditional on the $3 \times 1$ vector of priors $q_i(\mu, \varepsilon)$.

Next we model a player's ability to choose an optimal strategy given his prior by supposing $y_{ij}$ is computed with error. Let $\eta_{ij}$ denote the additive computation error for game $i$ and strategy $j$. A player chooses the strategy corresponding to the largest component of the $3 \times 1$ vector $y_i + \eta_i$. For tractibility, we assume each $\eta_{ij}$ is an independent and identically distributed Weibull noise. Hence, the choice probabilities have a conditional logit form (McFadden, 1974),

$$P_{ij}(\gamma, \mu, \varepsilon) \equiv \frac{\exp[\gamma y_{ij}(\mu, \varepsilon)]}{\sum_k \exp[\gamma y_{ik}(\mu, \varepsilon)]}, \tag{3}$$

where the parameter $\gamma$ is the precision of the player's expected utility calculation. Equation (3) specifies a three-parameter family of probabilistic choice functions.

If we had data on a large number of games, then we could estimate $(\gamma, \mu, \varepsilon)$ for each participant and then observe how the estimates are distributed in the participant population. However, our data on 12 games are hardly adequate for this purpose. Data contain only so much information, and it would be naive to think that point estimates of $48 \times 3 = 144$ parameters are conveying reliable detailed information about the distribution.

An alternative approach would be to estimate a non-parametric population distribution over $(\gamma, \mu, \varepsilon)$ space, using a grid with $h$ divisions per dimension. However, since $h^3$ parameters must be estimated, this approach is impractical for any reasonable $h$. On the other hand, we can use our theory to specify a non-uniform grid with a practical number of regions, hence a practical number of parameters.

We divide the $(\gamma, \mu, \varepsilon)$ space into five regions corresponding to the archetypes of our theory as follows:

(0) If $\gamma = 0$, then we have a level-0 type—uniform play.

(1) If $\gamma > 0$, $\mu = 0$, and $\varepsilon = 1$, then we have a level-1 type—a (perhaps imprecise) best response to the uniform distribution.

(2) If $\gamma > 0$, $\mu > 0$, and $\varepsilon = 1$, then we have a level-2 type—a (perhaps imprecise) best response to a (perhaps imprecise) best response to the uniform distribution.

(3) If $\gamma > 0$ and $\varepsilon = 0$, then we have a naive Nash-type—a (perhaps imprecise) best response to the Nash equilibrium prior.

(4) If $\gamma > 0.1$ and $\varepsilon \in (0, 1)$, then we have a "worldly" player who chooses a best response (perhaps imprecisely) to a prior based on a belief that some players are level-0, level-1, and naive Nash types.

Based on the payoff matrices in our experiment, a value of $\gamma$ less than 0.1 yields level-1 choice probabilities that are not significantly different from the uniform distribution. Therefore, for identification purposes we require $\gamma \geq 0.1$ for all types except level-0, and $\mu \geq 0.1$ for level-2 types. Similar identification considerations lead to the conditions that $\varepsilon \in (0.1, 0.9)$ for the worldly type.

Note that the fifth region covers most of the interior of $(\gamma, \mu, \varepsilon)$ space. Since the model forces all behavior into the three parameter family of probabilistic choice functions defined by Eq. (3), it will not be surprising to find some of the data falling in this region. Our label of "worldly" for this region suggests an interpretation of the underlying behavior; namely, worldly types understand the Nash equilibrium concept but (quite correctly) do not believe everyone else does; instead, this type believes the non-Nash players are level-0 and level-1 types (the proportion being captured by $\mu \geq 0$). A less flattering interpretation is as a "residual" category that captures behavior other than level-0, level-1, level-2, and naive Nash. We will address this issue of interpretation again in Sections 5.3, 5.4, and 7.

Corresponding to these archetypes, let $(\gamma_l, \mu_l, \varepsilon_l)$ denote the parameters for a level-$l$ type ($l = 1, 2, 3, 4$), where $l = 3$ denotes the naive Nash type and $l = 4$ denotes the worldly type. Note that for level-1, we restrict $\mu_1 = 0$ and $\varepsilon_1 = 1$, so only $\gamma_1$ is to be estimated. For level-2, we restrict $\varepsilon_2 = 1$, so $\gamma_2$ and $\mu_2$ are to be estimated. For level-3, we restrict $\varepsilon = 0$, so only $\gamma_3$ is relevant and to be estimated. For level-4, all three parameters are to be estimated. Thus, in all there are seven parameters to be estimated for the boundedly rational types.

### 2.2. *The Rational Expectations Type*

Let $\alpha_l$ denote the proportion of the population that is level-$l$ ($l = 0, 1, 2, 3, 4$), where $\alpha_l \geq 0$ and $\sum_{l=0}^{4} \alpha_l \leq 1$. For notational convenience, we will let $l = 5$ denote the RE type and $\alpha_5 \equiv 1 - \sum_{l=0}^{4} \alpha_l$ denote the proportion of the population consisting of RE types.

Let $p_{ij}^r$ denote the probability that an RE type chooses strategy $j$ in game $i$. Then, the prior of an RE type is

$$q_{ij}^r(p^r) \equiv \alpha_0 P_0 + \sum_{l=1}^{4} \alpha_l P_{ij}(\gamma_l, \mu_l, \varepsilon_l) + \alpha_5 p_{ij}^r. \tag{4}$$

The expected payoff is

$$y_{ij}^r(p^r) \equiv U_{ij} q_i^r(p^r). \tag{5}$$

Assuming independent and identically distributed Weibull computational errors, the probability that an RE type chooses strategy $j$ in game $i$ is

$$p_{ij}^r(\gamma_5) \equiv \frac{\exp[\gamma_5 y_{ij}^r(p^r)]}{\sum_k \exp[\gamma_5 y_{ik}(p^r)]}, \tag{6}$$

where $\gamma_5$ is the precision of the RE type's expected utility calculation. Equation (6), for $j \in \{1, 2, 3\}$, implicitly defines a fixed point, which we denote by $R_i(\gamma_5)$.

The fixed play of the boundedly rational types can be factored out of expected payoffs, leaving a reduced game among the RE types. $R_i$ is a logistic (or quantal-response) equilibrium of that reduced game (see McKelvey and Palfrey, 1994). Note that as the precision $\gamma_5 \to \infty$, $R_i$ approaches a perfect Nash equilibrium of the reduced game. Conversely, as $\gamma_5 \to 0$, $R_i$ approaches the uniform distribution. In order to statistically identify RE-type behavior from random level-0 type behavior, we initially require $\gamma_5 \geq 0.1$.

The task of solving for $R_i$ is non-trivial. Gradient-based algorithms were unsuccessful. Instead we took the Newton–Raphson rootfinder approach which utilized the matrix of partial derivatives of the right-hand side of Eq. (6). With an initial guess equal to the value of the right-hand side of Eq. (6) with $p^r = p^{\text{NE}}$, we obtained rapid convergence (often after only one iteration). While we cannot be sure that these reduced games have unique symmetric logistic equilibria for all possible parameter values, for parameter values near the maximum likelihood estimates the reduced games all have unique strictly dominant strategies.

### 2.3. *Integrating the Types into a Multimodal Model*

Let $s(h, i) \in \{1, 2, 3\}$ denote the strategy chosen by participant $h$ in game $i$, and let $s^h \equiv \{s(h, i), i = 1, \ldots, 12\}$ denote the joint choices of participant $h$. *Assuming that a participant's type is fixed for all games,* the probability of participant $h$'s joint choices conditional on being level-$l$ type is given by

$$P_l^h(\gamma_l, \mu_l, \varepsilon_l) \equiv \prod_i P_{is(h,i)}(\gamma_l, \mu_l, \varepsilon_l), \qquad \text{for } l = 1, 2, 3, 4,$$

and (7)

$$P_5^h(\gamma_5) \equiv \prod_i R_{is(h,i)}(\gamma_5).$$

The actual population of players can consist of all six types. We let $\beta$ denote the complete vector of $\alpha$, $\gamma$, $\mu$, and $\varepsilon$ parameters for all types. Then, the *ex ante* likelihood of participant $h$'s joint choices is given by

$$L(s^h|\beta) \equiv \alpha_0 P_0 + \sum_{l=1}^{4} \alpha_l P_l^h(\gamma_l, \mu_l, \varepsilon_l) + \alpha_5 P_5^h(\gamma_5), \tag{8}$$

with $\sum_{l=0}^{5} \alpha_l = 1$, and hence the log-likelihood of the observed sample is given by

$$\mathcal{L} \equiv \sum_h \log[L(s^h | \beta)]. \tag{9}$$

## 3. The Experimental Design

We designed an experiment to identify the parameters of the econometric specification, to control for idiosyncratic aspects of any one game, and to address the question of whether a player behaves as one type over a wide class of games. While it might be possible to select one game that could potentially identify the hypothesized types, the complexity of that game may inhibit the common knowledge understanding we want to induce. Instead, we selected a number of symmetric 3 × 3 games because they are relatively simple to understand and yet rich enough to permit identification from the observed behavior over all the games.[7] In addition, having data on a number of games allows us to investigate the predictive ability of our model by estimating the model on a subset of games and using these estimates to predict the behavior for the other games.

Twelve symmetric 3 × 3 games were selected with a variety of characteristics: three were strict dominance solvable [1, 5, 12], two were weak dominance solvable [3, 9], and three had unique mixed-strategy NE [4, 7, 11], while the remaining nine had unique pure-strategy symmetric NE. In four games [2, 6, 8, 10], the unique pure-strategy symmetric NE, $b(P_0)$, and $b^2(P_0)$ are distinct. The payoff matrices for the row player are presented in Table I; the transposes of these matrices give the payoffs for the column player.

A participant played each game exactly once and chose a single pure strategy for each game (always as a row player). Then each participant's choices were matched with every other participant. Each participant was essentially "playing against the field," i.e., against the empirical distribution of choices made by all other participants.[8] For each game, partici-

[7] The inability of many experimental designs to discriminate between alternative behaviors because of "flat likelihood" functions has been widely observed, e.g. Holt (1993). El-Gamal and Palfrey (1995) suggest a Bayesian design procedure, but the dimensionality of our model renders their method intractable. However, we employed the spirit of their criteria in our heuristic selection process.

[8] While pure theory would hold that playing against a single opponent randomly selected from the population is equivalent to playing against the field, we felt that the latter protocol would reinforce the inappropriateness of asymmetric equilibria. Friedman (1993) finds very little difference between the two protocols in learning-by-doing experiments; if anything the limiting behavior is slightly more Nash-like when playing against the field.

TABLE I

GAMES USED IN EXPERIMENT

| Game | | "T" | "M" | "B" | Game | | "T" | "M" | "B" |
|---|---|---|---|---|---|---|---|---|---|
| | | 7 | 40 | 1 | | | 21 | 17 | 10 |
| | T | 25 | 30 | 100 | | T | 30 | 100 | 50 |
| 1 | M | 40 | 45 | 65 | 7 | M | 40 | 0 | 90 |
| | B | 31 | 0 | 40 | | B | 50 | 75 | 29 |
| | | 30 | 12 | 6 | | | 12 | 12 | 24 |
| | T | 75 | 40 | 45 | | T | 0 | 60 | 50 |
| 2 | M | 70 | 15 | 100 | 8 | M | 100 | 20 | 50 |
| | B | 70 | 60 | 0 | | B | 50 | 40 | 52 |
| | | 5 | 16 | 27 | | | 26 | 1 | 21 |
| | T | 75 | 0 | 45 | | T | 40 | 100 | 65 |
| 3 | M | 80 | 35 | 45 | 9 | M | 33 | 25 | 65 |
| | B | 100 | 35 | 41 | | B | 80 | 0 | 65 |
| | | 26 | 15 | 7 | | | 39 | 3 | 6 |
| | T | 30 | 50 | 100 | | T | 45 | 50 | 21 |
| 4 | M | 40 | 45 | 10 | 10 | M | 41 | 0 | 40 |
| | B | 35 | 60 | 0 | | B | 40 | 100 | 0 |
| | | 14 | 3 | 31 | | | 13 | 4 | 31 |
| | T | 10 | 100 | 40 | | T | 30 | 100 | 22 |
| 5 | M | 0 | 70 | 50 | 11 | M | 35 | 0 | 45 |
| | B | 20 | 50 | 60 | | B | 51 | 50 | 20 |
| | | 11 | 20 | 17 | | | 26 | 3 | 19 |
| | T | 25 | 30 | 100 | | T | 40 | 15 | 70 |
| 6 | M | 60 | 31 | 51 | 12 | M | 22 | 80 | 0 |
| | B | 95 | 30 | 0 | | B | 30 | 100 | 55 |

pants' "token earnings" were computed, and this number gave the percentage chance of winning $2.00 for that game. A random number uniformly distributed on [00.0, 99.9] was generated by the throw of 3 color-coded 10-sided die. The player won $2.00 if and only if the his/her token earnings exceeded the random number. Actual payments were made immediately following the session. This method of computing token earnings and random money winnings was explained fully to all participants. The instructions (see the Appendix) were similar to Cooper *et al.* (1990). Actual monetary winnings ranged from $4.00 to $18.00, with an average of $11.63, and the whole experimental session took about 75 min.

The experiment was conducted in three sessions over a two-week period. The participants were fourth and fifth year undergraduate accounting and finance majors at the University of Texas. The first session had 14 participants, the second session had 22 participants, and the third session had 12 participants.

Following a training period (see the Appendix), each participant was given a 10 minute screening test (which everyone passed), designed to eliminate potential participants who did not understand the basics of the games and to instill common knowledge among all participants that everyone did understand these basics.

Each participant was given 36 min to complete their choices for the 12 games. Actual time used ranged from 15 to 36 min. Participants who finished early were required to sit quietly not doing anything until instructed to do so at the end of the 36 min.

We considered an alternative sequential protocol which would have allotted 3 min. for each game. Both protocols have advantages and disadvantages. Since some games are more difficult than others (e.g., games with no pure-strategy NE versus games that are strict-dominance solvable), our protocol allows participants to allot their time accordingly, while the alternative setup would induce more noise in the choices due to arbitrary time constraints. Also, some participants' thought processes may evolve during the experiment due to the variety of games in the experiment and the different experiences of thinking about these games. (Recall that they receive no feedback about other participants' choices, so we are describing a purely internal mental process.) Our setup will allow such participants to reconsider choices, thereby increasing the likelihood that their choices result from a single "model of the world," in contrast to the alternative setup. One disadvantage of our setup is that we will not have a truly "one-shot" environment, so the participants' choices should be interpreted as conditional on being presented with all 12 games. On the other hand, in the alternative setup, the choices should be interpreted as conditional on the past sequence of games, so no two choices would have identical conditioning events. We believe the net advantages of our setup exceed those of the alternative.

## 4. The Experimental Data

The aggregate choice data are given in Table I; the total number of participants making a particular choice are the underlined numbers above the respective payoff matrix. The disaggregate data are presented in Table II, which is sorted by our posterior identification of the participants' type (discussed in Section 5.3). The ID numbers were arbitrarily assigned by session, with the first, second, and third sessions having ID numbers 1–14, 15–36, and 37–48 respectively. The choices by game are given along with the actual token earnings. The average token earnings were 46.32 (s.d. = 2.35). In addition, Table II also gives the number of deviations from

TABLE II

ACTUAL DATA SORTED BY POSTERIOR TYPE

| Part | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | Token Earn | DEV NE | DEV b1 | DEV b2 | DEV RE | DM |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Level-0:** | | | | | | | | | | | | | | | | | | |
| 20 | M | T | M | T | B | B | M | T | T | T | B | M | 41.63 | 4 | 10 | 6 | 8 | 1 |
| 21 | T | M | B | T | T | T | M | B | T | M | B | B | 43.45 | 8 | 4 | 10 | 9 | 0 |
| 24 | M | T | T | T | M | M | B | T | B | M | M | T | 47.00 | 4 | 11 | 6 | 8 | 2 |
| 27 | M | T | T | M | B | M | T | B | M | T | T | B | 45.12 | 3 | 9 | 9 | 8 | 2 |
| 29 | M | M | B | M | T | T | T | M | T | T | B | T | 47.37 | 6 | 5 | 8 | 6 | 0 |
| 35 | M | T | T | M | M | M | T | B | B | T | T | M | 44.00 | 3 | 10 | 9 | 8 | 3 |
| 37 | M | M | B | T | B | B | M | T | T | T | B | M | 41.40 | 6 | 8 | 7 | 8 | 1 |
| 40 | B | M | T | T | T | T | T | M | T | T | T | T | 43.59 | 7 | 4 | 11 | 5 | 2 |
| **Level-1:** | | | | | | | | | | | | | | | | | | |
| 14 | T | M | B | T | T | T | T | M | T | T | T | B | 46.20 | 8 | 1 | 12 | 6 | 0 |
| 22 | M | M | B | T | T | T | T | M | T | B | T | B | 47.38 | 8 | 1 | 11 | 4 | 0 |
| 23 | T | M | B | T | T | T | T | M | T | B | T | B | 45.90 | 9 | 0 | 12 | 5 | 0 |
| 25 | T | M | B | T | T | T | T | M | T | B | T | B | 45.90 | 9 | 0 | 12 | 5 | 0 |
| 26 | M | T | B | T | T | M | T | M | T | B | T | B | 46.53 | 6 | 3 | 11 | 6 | 0 |
| 28 | T | M | B | T | T | T | T | M | T | B | T | B | 45.90 | 9 | 0 | 12 | 5 | 0 |
| 31 | T | M | B | T | T | M | B | M | T | T | T | B | 45.08 | 7 | 3 | 11 | 8 | 0 |
| 34 | M | T | B | T | T | M | T | M | T | T | B | B | 47.08 | 5 | 5 | 10 | 8 | 0 |
| 36 | T | B | T | T | M | T | T | M | T | B | T | B | 44.93 | 9 | 3 | 11 | 6 | 2 |
| 47 | M | M | B | T | T | T | T | M | T | T | T | B | 47.61 | 7 | 2 | 11 | 5 | 0 |
| **Level-2:** | | | | | | | | | | | | | | | | | | |
| 48 | M | B | M | M | B | B | B | T | B | M | B | T | 46.14 | 4 | 12 | 0 | 7 | 0 |
| **Level-3 (Naive Nash Types):** | | | | | | | | | | | | | | | | | | |
| 4 | M | T | B | M | B | M | T | B | T | T | B | T | 47.19 | 2 | 9 | 7 | 8 | 0 |
| 5 | M | T | B | M | B | M | T | B | T | T | B | T | 47.19 | 2 | 9 | 7 | 8 | 0 |
| 6 | M | T | M | M | B | M | B | B | B | T | B | T | 48.46 | 0 | 12 | 4 | 7 | 0 |
| 8 | M | T | B | M | B | M | B | B | T | T | B | T | 46.72 | 2 | 10 | 6 | 9 | 0 |
| 32 | M | T | B | M | B | M | T | B | T | T | B | T | 47.19 | 2 | 9 | 7 | 8 | 0 |
| 39 | M | T | B | T | B | M | M | B | T | T | M | T | 46.61 | 2 | 9 | 9 | 8 | 0 |
| 42 | M | T | M | M | B | M | M | B | B | T | B | T | 46.85 | 0 | 12 | 5 | 7 | 0 |
| 12 | M | T | M | T | B | M | M | B | B | T | B | T | 47.77 | 0 | 11 | 6 | 6 | 0 |
| 16 | M | T | M | T | B | M | B | B | B | T | B | T | 49.27 | 0 | 11 | 5 | 6 | 0 |
| 46 | M | T | M | T | B | M | T | B | B | T | B | T | 49.65 | 0 | 10 | 6 | 5 | 0 |

"ideal" archetypal behavior: specifically, from naive Nash behavior, $b(P_0)$, and $b^2(P_0)$. In addition, we also give the number of dominated strategies chosen. The "DEV RE" column will be explained in Section 5.3.

Note that 16 participants had 2 or fewer deviations from NE behavior,

TABLE II—*Continued*

| Part | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | Token Earn | DEV NE | DEV b1 | DEV b2 | DEV RE | DM |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Level-4 (Worldly Types):** | | | | | | | | | | | | | | | | | | |
| 1 | M | T | M | T | B | M | T | B | T | T | B | B | 47.49 | 2 | 8 | 8 | 7 | 0 |
| 2 | M | T | B | T | B | B | M | B | B | T | B | B | 46.23 | 3 | 9 | 7 | 8 | 0 |
| 3 | M | T | B | T | B | B | B | B | B | T | B | B | 47.87 | 3 | 9 | 6 | 8 | 0 |
| 7 | M | T | M | M | T | M | B | B | B | T | B | T | 47.46 | 1 | 11 | 5 | 8 | 0 |
| 9 | M | T | M | M | B | B | M | B | B | T | B | T | 45.98 | 1 | 12 | 4 | 7 | 0 |
| 10 | M | B | B | B | B | T | M | T | B | T | B | T | 46.88 | 4 | 10 | 5 | 7 | 0 |
| 11 | M | T | M | T | B | B | M | B | T | T | B | B | 44.61 | 3 | 9 | 7 | 8 | 0 |
| 13 | M | B | M | B | B | B | M | T | B | T | M | T | 45.64 | 3 | 12 | 4 | 7 | 0 |
| 15 | M | B | M | T | B | B | T | T | B | T | B | T | 47.87 | 3 | 10 | 3 | 5 | 0 |
| 17 | M | B | B | B | B | B | M | B | B | T | B | T | 45.69 | 3 | 11 | 5 | 8 | 0 |
| 18 | M | T | B | T | B | B | T | B | T | T | B | T | 47.17 | 3 | 8 | 7 | 7 | 0 |
| 19 | M | T | B | T | B | B | M | T | B | T | B | T | 46.03 | 3 | 10 | 5 | 7 | 0 |
| 30 | M | T | M | M | B | M | B | B | T | T | B | B | 46.23 | 2 | 10 | 6 | 9 | 0 |
| 33 | M | T | M | M | B | B | M | T | B | T | T | T | 45.25 | 2 | 11 | 4 | 6 | 0 |
| 38 | M | T | B | B | B | B | B | B | B | T | B | B | 47.18 | 3 | 10 | 6 | 9 | 0 |
| 41 | M | T | B | B | B | B | M | T | T | T | M | B | 43.80 | 5 | 9 | 8 | 10 | 0 |
| 43 | M | M | B | B | B | M | M | T | B | T | B | T | 45.96 | 3 | 10 | 6 | 7 | 0 |
| 44 | M | T | B | M | T | B | T | B | T | T | B | T | 45.16 | 4 | 8 | 7 | 9 | 0 |
| 45 | M | T | M | B | B | B | M | T | B | T | B | T | 45.31 | 2 | 12 | 4 | 7 | 0 |
| NE: | M | T | M | A | B | M | A | B | B | T | A | T | | 0 | 9 | 4 | 4 | |
| b1: | T | M | B | T | T | T | T | M | T | B | T | B | | 9 | 0 | 12 | 5 | |
| b2: | M | B | M | M | B | B | B | T | B | M | B | T | | 4 | 12 | 0 | 7 | |
| RE: | M | M | M | T | B | T | T | M | B | B | T | T | | 4 | 5 | 7 | 0 | |
| DM: | B | | T | | M | | | | M | | | M | | | | | | |

and 6 participants had 2 or fewer deviations from precise level-1 behavior,[9] while only one participant had 2 or fewer deviations from precise level-2 behavior. Thus, superficially the data appear to exhibit distinct patterns of behavior.

Directly below the individual choices, we give the choices of the ideal archetypes: naive Nash behavior, $b(P_0)$, $b^2(P_0)$, and the dominated strategies (DM). In the NE row, an entry of "A" denotes that all three strategies are in the support of the NE. The "RE" row will be explained in Section

[9] Since $b(P_0)$ coincides with the row with a payoff of 100 in all 12 games, an alternative interpretation of the underlying behavior is "maximax." In Stahl and Wilson (1994), the experimental design included several games in which level-1 and maximax behavior differed, and we found that a significant portion of participants exhibited distinctly level-1 behavior. This suggests that it is unlikely that all players we label as level-1 in the present data are instead maximax players.

5.3. The bottom, right-hand portion of Table II illustrates the relationship between the ideal archetypal behaviors on these 12 games. In 9 out of 12 games the Nash equilibrium choice is different from $b(P_0)$, and in 4 out of 12 games it is different from $b^2(P_0)$. In all 12 games $b(P_0)$ differs from $b^2(P_0)$.

While 6 out of 48 participants (12.5%) chose at least one strictly dominated strategy, only 7 choices were strictly dominated strategies in the three games [1, 5, 12] with strictly dominated strategies; thus, $7/(3 \times 48) = 4.86\%$ of the choices in these games were strictly dominated. Seven participants chose at least one strictly or weakly dominated strategy, bringing the total of dominated choices to 13 in the five relevant games [1, 3, 5, 9, 12]; thus, $13/(5 \times 48) = 5.42\%$ of the choices in these games were dominated.

Participants did not solve games by iterative elimination of strictly dominated pure strategies. In games 1, 5, and 12, this procedure leads to a unique choice (M, T, and T respectively), but 7, 17, and 22 participants (respectively) did not choose the iteratively undominated strategy.

Similarly, participants did not solve games by iterative elimination of weakly dominated strategies. In addition to the above games, games 3 and 9 possess a unique "iteratively admissible" strategy ($M$ and $B$ respectively); however, 32 and 27 participants (respectively) failed to choose the iteratively admissible strategy.

This is conclusive evidence that while most participants avoided strategies dominated with respect to all three strategies, their model of other players did not incorporate this behavior.

In the nine games with a unique pure-strategy symmetric NE, 42.8% of the responses differed from the unique NE. In two of the games with a unique mixed-strategy NE (4 and 7), the empirical distribution differs from the NE distribution at the 5% significance level. In game 11, the empirical distribution is insignificantly different from the mixed NE at the 5, 10, and 15% levels.

## 5. Statistical Analysis

The model represented by Eqs. (8), (9) constitutes a finite mixture model. Log-likelihood functions were maximized using the simplex method of Nelder and Mead (1965), using a variety of starting values to increase our confidence that a global maximum was achieved. The simplex method requires only function evaluations. Although it is not very efficient in terms of the number of function evaluations required, the method is easier to implement than other algorithms sometimes used with mixture models such as the EM-algorithm.

Nonparametric confidence intervals for parameter estimates were estimated using the bootstrap percentile method described by Efron (1982, Chap. 10). While it is possible to obtain conventional standard error estimates by evaluating the information matrix derived from the log-likelihood in Eqs. (8), (9), interpretation of $t$-ratios obtained from these estimates is problematic due to the proximity of the parameter estimates to parameter-space boundaries and the non-normality of the underlying distributions (as evidenced by our bootstrap results). The bootstrap procedure is particularly advantageous in our setting because it allows us to estimate the posterior probabilities that each player is of a particular type using a semi-parametric Bayesian procedure as discussed below.

The bootstrap method is based on the notion of replicating error processes by resampling estimated residuals. Since our model is a mixture of discrete choice models, residual terms are not explicitly estimated, and so the simulation step requires some modification. We first maximize the log-likelihood, Eq. (9), using the actual dataset to obtain a vector of parameter estimates $\hat{\beta}$. Then, to generate pseudo-data $s_*$ for each player, a uniform [0, 1] pseudo-random deviate is generated via the multiplicative congruential method and compared to the estimates $\hat{\alpha}_l$, $l = 0, \ldots, 5$, to determine player type. Next, uniform [0, 1] pseudo-random deviates are generated to determine a choice by the player on each of the 12 games in the experiment, using the estimated parameter values and Eqs. (3) and (6). For example, for a level-0 player a uniform [0, 1] pseudo-random deviate $v$ is generated: if $v \leq 0.\bar{3}$ then the choice for this particular game is recorded as 1; if $0.\bar{3} < v \leq 0.\bar{6}$ then the choice for this particular game is recorded as 2; and if $0.\bar{6} < v \leq 1$ then the choice is recorded as 3. For other player types, the process is similar, except that the intervals are determined by computing the discrete probability functions using the original parameter estimates in $\hat{\beta}$.

Once a $48 \times 12$ matrix of pseudo-choices $s_*$ has been simulated, the model is reestimated using these pseudo-data to obtain a bootstrap estimate $\hat{\beta}^*$. Then the process is repeated a large number of times to produce $M$ estimates, $\{\hat{\beta}^*(m)\}_{m=1}^{M}$. The bootstrap estimates $\hat{\beta}^*$ approximate the sampling distribution of the original estimator, $\hat{\beta}$. Let $\hat{\beta}_j$ and $\hat{\beta}_j^*(m)$ denote the $j$th elements of $\hat{\beta}$ and $\hat{\beta}^*(m)$, respectively. Nonparametric confidence intervals for $\hat{\beta}_j$ are obtained by sorting $\{\hat{\beta}_j^*(m)\}_{m=1}^{M}$ by algebraic value and then deleting the appropriate number of values from each end of the sorted array. If 95% confidence intervals are desired, then $0.025 \times M$ values would be deleted from each end of the sorted array; the new endpoints give the confidence interval. In the results reported below, we choose $M = 1000$ to ensure adequate coverage.

The bootstrap estimates $\hat{\beta}^*$ are also used to compute standard error estimates which are reported below. The bootstrap standard error estimates for $\hat{\beta}_j$ are obtained by computing the sample standard deviation of

$\{\hat{\beta}_j^*(m)\}_{m=1}^M$. As noted earlier, the interpretation of standard errors is problematic due to the non-normality of the underlying distributions.

### 5.1. *Estimation of the Mixture Model with RE Types*

We maximized the log-likelihood function in Eq. (9) for the entire sample of 48 participants, yielding a maximized value of $-442.390$. We found $\hat{\alpha}_5 = 0.03$, a bootstrapped confidence interval for $\alpha_5$ that included zero, and $\hat{\gamma}_5 = 0.1$, which is the minimum value initially imposed as an identification restriction. We first address the question of whether $\alpha_5$ is significantly different from zero, and secondly whether our identification restriction was adequate.

We reestimated the model while restricting $\alpha_5 = 0$, which yielded a log-likelihood of $-442.727$. The corresponding likelihood-ratio statistic, $\hat{\chi}^2 = -2(\mathcal{L}_{\text{restricted}} - \mathcal{L}_{\text{free}}) = 0.674$, would have a $p$-value of 0.714 if $\hat{\chi}^2$ had the usual chi-square distribution with 2 degrees of freedom. Unfortunately, however, the likelihood-ratio statistic has unknown distribution under the null hypothesis since the null value of $\alpha_5$ is on the boundary of the parameter space (see Everitt and Hand, 1981, and Titterington *et al.*, 1985, for discussions of this problem in the context of finite mixture models). Conventional Wald and Lagrange multiplier tests are also invalid at the edge of the parameter space.

To circumvent this problem, we used the bootstrap procedure to approximate the sampling distribution of the likelihood-ratio statistic. The choice data were simulated as outlined above for the bootstrap procedure (except $\hat{\beta}_{\text{restricted}}$ was used to simulate the pseudo-data since the null hypothesis is $\alpha_5 = 0$). Both the restricted and unrestricted models were estimated on these pseudo-data yielding log-likelihood values $\mathcal{L}^*_{\text{restricted}}$ and $\mathcal{L}^*_{\text{free}}$, respectively. These values were then used to compute a bootstrap estimate of the likelihood-ratio statistic: $\hat{\chi}^2_* = -2(\mathcal{L}^*_{\text{restricted}} - \mathcal{L}^*_{\text{free}})$. The entire process was repeated 1000 times to produce bootstrap estimates $\{\hat{\chi}^2_*(m)\}_{m=1}^{1000}$. Since these values approximate the sampling distribution of the original likelihood-ratio statistic, it is straightforward to determine the significance of the original likelihood-ratio statistic by first sorting the values $\{\hat{\chi}^2_*(m)\}_{m=1}^{1000}$ and then determining the percentile of the original statistic. The value 0.674 has a $p$-value of 0.165, clearly indicating that we cannot reject the hypothesis that there are no RE types in the sample population.[10]

[10] Our bootstrap procedure to test $\alpha_5 = 0$ is methodologically identical to the Monte Carlo approach employed by Aitken *et al.* (1981), who in effect used only 19 bootstrap replications. Hall (1986) provides theoretical results which show that using a small number of replications in the bootstrap may increase the probability of type-II errors. Several authors have suggested using at least 100 replications for testing null hypotheses regarding parameter values and as many as 1000 replications for constructing confidence intervals. Note that merely examining the confidence interval ignores the variation in the other parameters; thus, the bootstrapped likelihood-ratio statistic provides a more powerful test.

TABLE III

PARAMETER ESTIMATES FOR UNIMODAL MODELS

| | Level-1 | Level-2 | Level-3 | Level-4 | Level-5 |
|---|---|---|---|---|---|
| $\mu$ | — | 0.045 | — | 0.040 | — |
| $\epsilon$ | — | — | — | 0.603 | — |
| $\gamma$ | 0.034 | 0.107 | 0.082 | 0.141 | 0.064 |
| $\mathcal{L}$ | −605.29 | −558.79 | −576.10 | −544.27 | −556.32 |

We also examined the choice probabilities, $R_i$, and the payoff matrices for the reduced game between RE types, evaluated at the maximum likelihood estimates. The former are quite diffuse despite the fact that every reduced game had a strictly dominant strategy. Since it would be more reasonable for RE types to be more precise about recognizing dominant strategies, we suspect that our restriction that $\gamma \geq 0.1$ was not an adequate identification restriction.

To investigate this issue further, we computed a goodness-of-fit statistic, $\lambda = \Sigma_i \Sigma_j (R_{ij}(\hat{\beta}) - P_0)^2/P_0$, to compare these choice probabilities with the uniform distribution. We found that $\lambda = 3.4$, which with 2 degrees of freedom is insignificant. In other words, with a sample of 12 observations, it would not be possible to identify a participant as an RE type or a level-0 type at any acceptable degree of confidence. When we restricted $\gamma_5$ to be large enough for identification at the 5% level ($\geq 0.175$), the maximum likelihood estimate of $\alpha_5$ was driven down to 0.005. It is also noteworthy that the only boundedly rational type for which $\hat{\alpha}_l$ changed was the level-0 type. Thus, it appears that our original identification restriction was inadequate and allowed one or two level-0 types to appear to be very error-prone RE types.[11]

Since (i) our bootstrapped test did not reject the null hypothesis of $\alpha_5 = 0$, and (ii) $\hat{\alpha}_5 = 0$ when we restricted $\gamma_5$ to be sufficient for identification, we conclude that there are no RE types in our sample. However, since our experimental design focused exclusively on initial choices by "inexperienced" participants, we obviously cannot rule out the possibility that RE behavior might be learned with experience.

As a benchmark for later comparisons, we also estimated five unimodal models with only one type of player: level-$l$ for $l = 1, \ldots, 5$. The maximum likelihood point estimates and the log-likelihoods are reported in Table III. Note that all the unimodal precision estimates (i.e., $\hat{\gamma}_l$) are small,

[11] See also footnote 17. This goodness-of-fit criteria turns out to be equivalent to the requirement that the Neymann–Pearson decision rule for the 95% confidence level have a 95% probability of being decisive.

TABLE IV

PARAMETER ESTIMATES AND CONFIDENCE INTERVALS FOR MIXTURE MODEL WITHOUT RE TYPES

| | Estimate | Std. Dev. | 95 percent conf. int. | |
|---|---|---|---|---|
| $\gamma_1$ | 0.2177 | 0.0425 | 0.1621 | 0.3055 |
| $\mu_2$ | 0.4611 | 0.0616 | 0.2014 | 0.8567 |
| | | | [0.2360 | 0.8567] |
| $\gamma_2$ | 3.0785 | 0.5743 | 1.9029 | 4.9672 |
| | | | [2.5631 | 5.0000] |
| $\gamma_3$ | 4.9933 | 0.9357 | 1.9964 | 5.0000 |
| $\mu_4$ | 0.0624 | 0.0063 | 0.0527 | 0.0774 |
| $\epsilon_4$ | 0.4411 | 0.0773 | 0.2983 | 0.5882 |
| $\gamma_4$ | 0.3326 | 0.0549 | 0.2433 | 0.4591 |
| $\alpha_0$ | 0.1749 | 0.0587 | 0.0675 | 0.3047 |
| $\alpha_1$ | 0.2072 | 0.0575 | 0.1041 | 0.3298 |
| $\alpha_2$ | 0.0207 | 0.0202 | 0.0000 | 0.0625 |
| $\alpha_3$ | 0.1666 | 0.0602 | 0.0600 | 0.2957 |
| $\alpha_4$ | 0.4306 | 0.0782 | 0.2810 | 0.5723 |
| $\mathcal{L}$ | -442.727 | | | |

making the predicted unimodal choice probabilities not much different from random noise. As a comparison, the level-0 model yields a log-likelihood of −632.801, while the log-likelihood for the mixture model without the RE types estimated below is −442.727. The likelihood ratio statistics comparing the unimodal models to the mixture model have chi-square $p$-values less than $10^{-42}$, indicating that the mixture model is vastly superior.[12]

## 5.2. *Estimation of the Mixture Model without RE Types*

Having concluded that there are no RE types in our sample, we exclude the RE type from the subsequent statistical analysis by restricting $\alpha_5 = 0$.[13] Estimation of the model without the RE types using the entire sample of 48 participants yielded the results shown in Table IV. To test whether

[12] As noted above, the likelihood ratio statistics do not have the familiar chi-square distribution; however, our experience with bootstrapping likelihood-ratio statistics suggests that these would have $p$-values far below the values required to reject the underlying null hypotheses.

[13] Alternatively, we could allow $\alpha_5$ to vary; however, given that no RE types exist, this would result in inefficient estimates for the remaining parameters.

analyses of the entire sample are valid, or if separate analyses by session should be conducted, we estimated the model for each session sample. The sum of the three log-likelihoods was −428.824, compared to −442.727 for the aggregated model. Since the aggregate model involves 22 restrictions relative to the case where parameters are allowed to vary across sessions, the likelihood-ratio statistic is distributed chi-square with 22 degrees of freedom. The computed value of 27.806 has a $p$-value of 0.182, and so we fail to reject the restrictions implied by the aggregate model.[14] Henceforth, we focus on analyses of the entire sample.

Note from Table IV that the bootstrapped 95% confidence intervals of the mixture parameters, $\alpha_0$, $\alpha_1$, $\alpha_3$, and $\alpha_4$, are strictly positive; only the confidence interval for $\alpha_2$ includes zero. This suggests that each type, except possibly the level-2 type, is present in the population. On the other hand, with $\hat{\alpha}_2 = 0.0207$, in each pseudo-data set of 48 participants there is a 37% chance that no level-2 types are present, in which case the estimate $\hat{\alpha}_2^*$ is very likely to be zero. Hence, it is not surprising that the bootstrapped confidence interval has 0 as the lower bound.

Restricting $\alpha_2 = 0$ and reestimating the model yields a log-likelihood of −446.582, and hence a likelihood-ratio statistic of 7.711. Unfortunately, the likelihood-ratio statistic again has unknown distribution since the null hypothesis is at the boundary of the parameter space. Thus, we employ the bootstrap methodology described above. The value 7.711 turns out to be significant at the .012 confidence level. Thus, we reject the null hypothesis $\alpha_2 = 0$ and conclude that level-2 type players are present in the data. Moreover, Bayesian posterior procedure described below in Section 5.3 identifies Participant 48 as a level-2 type, and referring to Table II, we see that this participant's choices are exactly $b^2(P_0)$; hence we feel confident that there was at least one level-2 type in the sample population.

In bootstrapping the confidence intervals shown in Table IV, we found that 354 of the 1000 bootstrap replications produced bootstrap estimates $\hat{\alpha}_2^* = 0$. In such cases, the bootstrap estimates $(\hat{\mu}_2^*, \hat{\gamma}_2^*)$ are spurious; they can drift to arbitary values during the optimization process because the likelihood function is obviously insensitive to these parameters when $\hat{\alpha}_2^* = 0$. In order to eliminate this additional source of noise reflected in the confidence intervals for $\hat{\mu}_2$ and $\hat{\gamma}_2$ shown in Table IV, we used the 646 bootstrap estimates $\{\hat{\mu}_2^*(m), \hat{\gamma}_2^*(m) | \hat{\alpha}_2^*(m) > 0, m = 1, \ldots, 1000\}$ to compute *conditional* 95% confidence intervals, shown in brackets in Table IV.

[14] The same exercise for the full model with RE types yields the same conclusion. Since we are not at the boundary of the parameter space under the null in this test, the likelihood-ratio statistic has the usual chi-square distribution.

The upper bound of exactly 5.0 on the confidence intervals for $\gamma_2$ and $\gamma_3$ might look suspicious to the reader. Indeed, this reflects the upper bound we imposed in order to guarantee speedy convergence. The probabilistic choice functions for the types have the property that as the precision ($\gamma_l$) becomes large, one strategy in each game is chosen with probability indistinguishably different from one. By continuity, as $\gamma_l$ increases, the probabilistic choice function becomes insensitive to $\gamma_l$. Consequently, the bootstrapped distribution of precision parameters can have a large upper tail, *but almost all of these are behaviorally equivalent*. Through experimentation, we found that increasing the upper bound from 5 to 10 left the maximized log-likelihood value and all other parameter estimates unchanged to 4 significant digits; therefore, we settled on an upper bound of 5 for each $\gamma_l$, $l = 1, \ldots, 4$. On the other hand, there would be no statistically significant difference in our results if we were to have imposed an upper bound of 3, so values of $\gamma_l > 3$ should be considered statistically and behaviorally equivalent, and the reader should not be distracted by large upper confidence bounds for the $\gamma_l$.

The lower confidence bounds for the $\gamma_l$, $l = 1, \ldots, 4$, however, are vitally important for identification purposes as pointed out in Section 2. These all exceed the required minimum of 0.1. Conditioning on $\hat{\alpha}_2^* > 0$ increases the lower bound on the confidence interval for $\gamma_2$ from 1.9029 to 2.5631. The predicted probabilistic choice functions for each type are quite different from the uniform distribution and from each other, indicating successful identification.

While we have ruled out the presence of RE types in Section 5.1, we have yet to address the question of whether some subset of types could be deleted without a significant deterioration of the likelihood function. Since there are five archetypes, there are 30 distinct hypotheses, each involving some combination of the $\alpha_l$, $l = 0, \ldots, 4$, being restricted to zero.[15] For these hypotheses, we computed the likelihood-ratio statistic using the actual data. As noted earlier, the null hypothesis $\alpha_2 = 0$ produced a likelihood ratio statistic of 7.711. The bootstrap procedure employed above indicated that this value is significant at .012; if the statistic were assumed to be chi-square with 3 degrees of freedom, the corresponding $p$-value would be 0.053, suggesting that the actual distribution of the likelihood-ratio statistic has a thinner upper tail than the chi-square distribution. Among the remaining 29 hypotheses to be tested, the next-highest chi-square $p$-value was less than $10^{-5}$, and so we infer that the remaining null hypotheses would also be rejected by our bootstrap procedure.

[15] Five hypotheses result when either one or four of the $\alpha_l$ are restricted to zero; ten hypotheses result when either two or three of the $\alpha_l$ are restricted to zero. The hypothesis that all archetypes are absent is not pertinent since level-0 is the minimal null hypothesis.

### 5.3. *Computation of Posteriors*

Bayes' theorem can be used to derive the posterior probability that any participant $h$ is level-$l$ ($l = 0, 1, 2, 3, 4$), which we denote as $\alpha_l^h$. For notational convenience, let $s = \{s^h, h = 1, \ldots, 48\}$, the vector of all observations, let $L(s|\beta)$ denote the likelihood of the observed data $s$ conditional on the model parameters $\beta$, and let $f(\beta)$ denote the joint density of $\beta$. Then $\alpha_l^h \equiv L(\text{player } h \text{ is level-}l|s)$, and, by Bayes' theorem,

$$\alpha_l^h = \frac{L(\text{player } h \text{ is level-}l \text{ and } s)}{L(s)}. \tag{10}$$

Conditional probabilities can always be integrated to yield unconditional probabilities, and hence

$$L(s) = \int \cdots \int L(s|\beta) f(\beta)\, d\beta. \tag{11}$$

Recall that $L(s^h|\beta)$ denotes the joint probability of observations corresponding to player $h$. Then $L(s|\beta) = \Pi_{h=1}^{48} L(s^h|\beta)$, and

$$\begin{aligned} L[(\text{player } h \text{ is level-}l \text{ and } s)|\beta] &= \alpha_l P_l^h \prod_{j \neq h} L(s^j|\beta) \\ &= \alpha_l P_l^h L(s|\beta)/L(s^h|\beta), \end{aligned} \tag{12}$$

and so

$$\alpha_l^h = \int \cdots \int \frac{\alpha_l P_l^h L(s|\beta)}{L(s^h|\beta) L(s)} \cdot f(\beta)\, d\beta. \tag{13}$$

The bootstrap procedure used to construct confidence intervals provides an approximation to the sampling distribution of $\beta$ which is used to perform the integration in (11) and (13). The integral in (11) is approximated by first computing $L(s|\hat{\beta}^*(m))$ for each bootstrap replication $m = 1, \ldots, M$ and computing the mean $M^{-1} \Sigma_{m=1}^{M} L(s|\hat{\beta}^*(m))$. The integral in (13) is computed similarly, using the value for $L(s)$ obtained from (11).[16] Consequently, the computation of the $\alpha_l^h$ amounts to a semi-parametric Bayesian procedure; the procedure is semi-parametric since the functional form of $f(\beta)$ is not specified. The procedure is more general than a fully parametric procedure that would require specification of $f(\beta)$.

[16] Note that $s$ is the actual experimental data and not the pseudo data.

The computed posterior probabilities $\alpha_l^h$ are presented in Table V. It is remarkable that 38 of the 48 participants have a 90% or better probability of being one type. Of these, 6 are level-0, 9 are level-1, 1 is level-2, 5 are naive Nash, and 17 are worldly types. This finding suggests that indeed the population of human players is not homogeneous in their thought processes, but rather is composed of distinct types of strategic thinking.

Further, for 44 of the 48 participants, the likelihood ratio of the most likely type to the next most likely type exceeds 2.4. Of these, 8 are level-0, 9 are level-1, 1 is level-2, 7 are naive Nash, and 19 are worldly. This finding suggests that most of the participants can be characterized as employing one type of thought process (or model of other players) for all 12 games (see also Section 6). Of the remaining four participants, three are roughly equally likely to be naive Nash or worldly types, and the remaining one (Participant 36) is most likely a level-1 type but with lower precision.

The 29 participants identified above as level-3 or level-4 appear to be cognizant of the Nash equilibrium concept, but the remaining 19 (40%) do not. Also recall that 42.8% of the responses on games with pure-strategy NE were non-NE. It is noteworthy that $\hat{\varepsilon}_4$, the worldly type's estimate of non-Nash behavior, was 0.44. Of course, this similarity may be a mere coincidence. Moreover, only about 15% of the population were naive Nash types, which is much lower than $\hat{\varepsilon}_4$.

Table II displays the actual choice data sorted by posterior type. The next to last row, labeled "RE," gives the best response to the aggregate mixture model predictions [see Eq. (14) below], which is what a single, perfectly precise RE-type player would choose against this population. The last column shows the number of deviations from this "RE-dominant" strategy for each participant. The first group of eight participants in Table II are the level-0 types. Evidently, these participants' choices are not explained by any of the archetypal behaviors.[17] The next group of 10 are the level-1 types (including Participant 36); these exhibit the fewest deviations from $b(P_0)$. Next, the singleton group of level-2 types is displayed; note that Participant 48's behavior is exactly fit by the level-2 type. The next group of 10 consists of the seven participants definitely identified as naive Nash (level-3) types plus three more (the last three) who were identified as either naive Nash or worldly types. Notably, the choices involve very few deviations from NE responses. The last group of 19 participants are the worldly (level-4) types.

[17] From the analysis of the model with RE types (with $\gamma_5 = 0.1$), only Participant 29 had a posterior $\alpha_5^h > 0.5$. However, we can see from Table II that his actual behavior involves six deviations from the RE-dominant strategy. Indeed, it is clearly evident from Table II that none of the participants' behavior is best fit by the rational expectations hypothesis.

TABLE V

POSTERIOR ESTIMATES OF PLAYERS' TYPES

| $h$ | $\alpha_0^h$ | $\alpha_1^h$ | $\alpha_2^h$ | $\alpha_3^h$ | $\alpha_4^h$ |
|---|---|---|---|---|---|
| 1 | 0.033 | 0.025 | 0.0 | 0.0 | 0.942 |
| 2 | 0.003 | 0.0 | 0.0 | 0.0 | 0.997 |
| 3 | 0.003 | 0.0 | 0.0 | 0.0 | 0.997 |
| 4 | 0.0 | 0.0 | 0.0 | 0.941 | 0.059 |
| 5 | 0.0 | 0.0 | 0.0 | 0.941 | 0.059 |
| 6 | 0.0 | 0.0 | 0.0 | 0.708 | 0.292 |
| 7 | 0.030 | 0.0 | 0.0 | 0.0 | 0.970 |
| 8 | 0.0 | 0.0 | 0.0 | 0.945 | 0.055 |
| 9 | 0.001 | 0.0 | 0.0 | 0.027 | 0.972 |
| 10 | 0.006 | 0.0 | 0.0 | 0.0 | 0.994 |
| 11 | 0.030 | 0.0 | 0.0 | 0.0 | 0.970 |
| 12 | 0.0 | 0.0 | 0.0 | 0.578 | 0.422 |
| 13 | 0.008 | 0.0 | 0.0 | 0.0 | 0.992 |
| 14 | 0.0 | 1.0 | 0.0 | 0.0 | 0.0 |
| 15 | 0.001 | 0.0 | 0.0 | 0.0 | 0.999 |
| 16 | 0.0 | 0.0 | 0.0 | 0.561 | 0.439 |
| 17 | 0.002 | 0.0 | 0.0 | 0.0 | 0.998 |
| 18 | 0.002 | 0.0 | 0.0 | 0.083 | 0.915 |
| 19 | 0.0 | 0.0 | 0.0 | 0.0 | 1.0 |
| 20 | 0.970 | 0.0 | 0.0 | 0.0 | 0.030 |
| 21 | 0.709 | 0.290 | 0.0 | 0.0 | 0.001 |
| 22 | 0.0 | 1.0 | 0.0 | 0.0 | 0.0 |
| 23 | 0.0 | 1.0 | 0.0 | 0.0 | 0.0 |
| 24 | 0.990 | 0.0 | 0.0 | 0.0 | 0.010 |
| 25 | 0.0 | 1.0 | 0.0 | 0.0 | 0.0 |
| 26 | 0.0 | 1.0 | 0.0 | 0.0 | 0.0 |
| 27 | 0.999 | 0.0 | 0.0 | 0.0 | 0.001 |
| 28 | 0.0 | 1.0 | 0.0 | 0.0 | 0.0 |
| 29 | 0.802 | 0.039 | 0.0 | 0.0 | 0.159 |
| 30 | 0.067 | 0.0 | 0.0 | 0.0 | 0.933 |
| 31 | 0.002 | 0.998 | 0.0 | 0.0 | 0.0 |
| 32 | 0.0 | 0.0 | 0.0 | 0.941 | 0.059 |
| 33 | 0.002 | 0.0 | 0.0 | 0.0 | 0.998 |
| 34 | 0.025 | 0.959 | 0.0 | 0.0 | 0.016 |
| 35 | 1.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| 36 | 0.348 | 0.652 | 0.0 | 0.0 | 0.0 |
| 37 | 0.993 | 0.0 | 0.0 | 0.0 | 0.007 |
| 38 | 0.008 | 0.0 | 0.0 | 0.0 | 0.992 |
| 39 | 0.0 | 0.0 | 0.0 | 0.968 | 0.032 |
| 40 | 0.985 | 0.015 | 0.0 | 0.0 | 0.0 |
| 41 | 0.206 | 0.0 | 0.0 | 0.0 | 0.794 |
| 42 | 0.0 | 0.0 | 0.0 | 0.722 | 0.278 |
| 43 | 0.007 | 0.0 | 0.0 | 0.0 | 0.993 |
| 44 | 0.144 | 0.0 | 0.0 | 0.0 | 0.856 |
| 45 | 0.001 | 0.0 | 0.0 | 0.0 | 0.999 |
| 46 | 0.0 | 0.0 | 0.0 | 0.543 | 0.457 |
| 47 | 0.0 | 1.0 | 0.0 | 0.0 | 0.0 |
| 48 | 0.0 | 0.0 | 0.981 | 0.0 | 0.019 |

TABLE VI

ACTUAL AND PREDICTED PROBABILITIES

| 1 | 2 | 3 | 1 | 2 | 3 |
|---|---|---|---|---|---|
| **(a) Actual Choice Frequencies:** | | | **(b) Mixture Model Predictions:** | | |
| 0.146 | 0.833 | 0.021 | 0.186 | 0.756 | 0.059 |
| 0.625 | 0.250 | 0.125 | 0.513 | 0.280 | 0.208 |
| 0.104 | 0.333 | 0.563 | 0.061 | 0.409 | 0.530 |
| 0.542 | 0.313 | 0.146 | 0.556 | 0.254 | 0.190 |
| 0.292 | 0.063 | 0.646 | 0.224 | 0.097 | 0.679 |
| 0.229 | 0.417 | 0.354 | 0.268 | 0.433 | 0.298 |
| 0.438 | 0.354 | 0.208 | 0.445 | 0.251 | 0.304 |
| 0.250 | 0.250 | 0.500 | 0.259 | 0.312 | 0.429 |
| 0.542 | 0.021 | 0.438 | 0.361 | 0.120 | 0.519 |
| 0.813 | 0.063 | 0.125 | 0.503 | 0.218 | 0.280 |
| 0.271 | 0.083 | 0.646 | 0.416 | 0.189 | 0.395 |
| 0.542 | 0.063 | 0.396 | 0.633 | 0.059 | 0.308 |

### 5.4. *Goodness of Fit*

The estimated $\hat{\beta}$ can be used to predict the aggregate choices $n_{ij}$, for game $i$ and strategy $j$. Let $\pi_{ij}$ denote the predicted probability of choice $j$ in game $i$ by a randomly selected participant. Then

$$\pi_{ij} \equiv \sum_{l=0}^{4} \hat{\alpha}_l P_{ij}(\hat{\gamma}_l, \hat{\mu}_l, \hat{\varepsilon}_l), \tag{14}$$

and the predicted aggregate choices are $N\pi_{ij}$, where $N$ is the number of participants. Table VI shows the actual distribution of choices, $n/48$, and the predicted distribution $\pi$ for the mixture model. A casual perusal of Table VI for the mixture model reveals a surprisingly good fit except for possibly games 10 and 11. We emphasize that our estimation procedure optimized for the best fit of the individual choice data (a $48 \times 12$ matrix) and consequently did *not* optimize for the best fit of the aggregate choice data (a $12 \times 3$ matrix).

The statistic

$$\lambda \equiv \sum_{i} \sum_{j} \frac{(n_{ij} - N\pi_{ij})^2}{N\pi_{ij}} \tag{15}$$

measures goodness-of-fit and is distributed chi-square with 24 degrees of freedom (the number of strategies minus one times the number of games). For the entire sample of 48 participants and 12 games, $\lambda = 57.57$ which is significant at the 0.01% level. However, for the 10 games excluding games 10 and 11, $\lambda = 26.09$ which, with 20 degrees of freedom, is not significant at any commonly accepted level. Therefore, we feel that the estimated model does a good job of fitting the aggregate distribution of choices.

It is interesting to ask how the other types (the estimated types of the mixture model as well as the estimated unimodal types) do at predicting the aggregate choices of all 48 participants. The next best fit is achieved by the unimodal level-4 model (given in Table III), and $\lambda = 100.8$ which has a $p$-value of $2 \times 10^{-11}$. Thus, while the predicted aggregate choices differ from the actual aggregate data, our model does much better than the alternatives.

We also computed the predicted distribution $\pi$ for each type subgroup of Table II. All distributions are distinct from one another, confirming that our identification restrictions were adequate.

Since we have grouped participants by type in Table II it is also interesting to ask how the estimated choice probabilities by type fit the aggregate behavior of the associated group of participants. Equation (15), with $N$ adjusted appropriately, can be used to gauge the goodness-of-fit for each type subgroup. For the 8 participants identified above as level-0 types, $\pi$ is uniform and $N = 8$, so Eq. (15) yields $\lambda = 31.50$, which is not significant at any commonly accepted level. Similarly, for the 10 participants identified as level-1 types, $\lambda = 23.824$; and for the 7 participants identified as naive Nash types, $\lambda = 24.81$, neither of which is significant at any commonly accepted level. Further, adding the three participants who are either naive Nash or worldly to the naive Nash group, we still cannot reject the hypothesis that the group's choices were generated by the naive Nash model. Thus, the observed choices of the players are consistent with their identified type.

The aggregate choices of the 19 participants identified as worldly types, however, are significantly different from the model predictions ($\lambda = 72.64$). The greatest discrepancies appear in games 6, 9, 10, and 11; however, eliminating these games, there is still a statistically significant difference. Also, adding the three participants who are either naive Nash or worldly to the worldly group does not significantly improve the fit.

We can also ask how well the estimated model fits the $48 \times 12$ matrix of disaggregated data. There are $3^{12} = 531441$ possible behavioral patterns, of which we observe only 44; two patterns are observed 3 times each. The probability of each of the $3^{12}$ patterns predicted by the model is given by Eq. (8) evaluated at the maximum likelihood parameter estimates $\hat{\beta}$.

TABLE VII

EXPECTED PAYOFFS AGAINST MODEL BY TYPE

| Game | Level-0 | Level-1 | Level-2 | Level-3 | Level-4 | BR | AVG | S.D. |
|---|---|---|---|---|---|---|---|---|
| 1 | 28.8 | 38.1 | 45.2 | 45.2 | 45.1 | 45.2 | 41.3 | 6.1 |
| 2 | 57.5 | 60.5 | 52.7 | 59.0 | 57.3 | 60.9 | 58.0 | 2.7 |
| 3 | 37.9 | 42.2 | 43.0 | 42.6 | 42.6 | 43.0 | 41.9 | 1.8 |
| 4 | 39.6 | 48.3 | 35.6 | 39.6 | 42.4 | 48.4 | 42.3 | 4.7 |
| 5 | 43.3 | 41.1 | 50.1 | 50.1 | 49.3 | 50.1 | 47.3 | 3.7 |
| 6 | 44.3 | 47.5 | 38.5 | 44.7 | 42.6 | 49.5 | 44.5 | 3.5 |
| 7 | 49.6 | 53.0 | 49.9 | 49.6 | 49.8 | 53.7 | 50.9 | 1.7 |
| 8 | 47.2 | 52.8 | 40.2 | 47.7 | 45.6 | 53.6 | 47.8 | 4.5 |
| 9 | 57.1 | 60.2 | 62.6 | 57.1 | 62.2 | 62.6 | 60.3 | 2.4 |
| 10 | 37.7 | 41.4 | 31.8 | 39.4 | 37.3 | 39.4 | 37.8 | 3.0 |
| 11 | 37.0 | 39.9 | 38.6 | 37.0 | 37.9 | 40.1 | 38.4 | 1.2 |
| 12 | 36.1 | 41.8 | 47.8 | 47.8 | 47.1 | 47.8 | 44.7 | 4.4 |
| AVG | 43.0 | 47.2 | 44.7 | 46.6 | 46.6 | 49.5 | 46.3 | |
| S.D. | 8.3 | 7.5 | 8.2 | 6.5 | 7.0 | 7.0 | 6.8 | |

Call this $L_k$, and let $n_k$ equal the number of times the $k$th pattern is observed, $k = 1, \ldots, 3^{12}$. Then the goodness-of-fit statistic given in Eq. (15) adapted to this test becomes $\lambda \equiv \Sigma_k (n_k^2/48L_k) - 48 = 530157$, which is distributed chi-square with $3^{12}$ degrees of freedom and has a $p$-value of 0.894. Thus, we cannot reject the hypothesis that the actual disaggregated data were generated by the model.

## 5.5. *Comparative Performance of Types*

Using $\pi$ computed in Eq. (14), the expected payoff to strategy $j$ in game $i$ is given by $\hat{y}_{ij} \equiv \Sigma_k U_{ijk}\pi_{ik}$. Hence, using Eq. (3), the expected payoff to a level-$l$ type in game $i$ is

$$\mathscr{E}_{il} \equiv \sum_j P_{ij}(\gamma_l, \mu_l, \varepsilon_l)\hat{y}_{ij}. \tag{16}$$

We can also calculate the expected payoff of the best response (BR) to the estimated model. Table VII presents this 12 × 6 matrix of expected payoffs, as well as the column and row averages and standard deviations. Casual observation reveals very little difference between the columns, except perhaps for the level-0 column. The chi-square contingency statistic

$$\sum_i \sum_l [(6\mathscr{E}_{il} - \mathscr{E}_{i\cdot})^2/6\mathscr{E}_{i\cdot}] \sim \chi^2\,(60\text{ d.f.}), \tag{17}$$

where $\mathscr{E}_{i\cdot} \equiv \Sigma_l \mathscr{E}_{il}$, may be used to test the hypothesis that the six columns are identical. The computed value of the statistic is 21.324 which has a $p$-value of $1 - 10^{-6}$; thus, we cannot reject the null hypothesis that all levels perform equally well over these 12 games against the predicted population. Further, the only pairwise payoff differences that are significant at the 10% level for a one-tailed $t$-test are level-0 versus level-1 and level-0 versus BR.

Harrison (1989) argued that a "flat" payoff structure, such as revealed by Table VII, indicates that the experimental participants might not have had sufficient incentives for high-level processing and problem solving, and hence any less-than-fully-rational observed behavior may be an artifact of insufficient incentives. However, we should also examine the incentives in each game. The payoff difference between the best and the next-best response to a pure strategy averaged over all pure strategies was 23.56 (or \$0.47) per game, and the payoff difference between the best and the worst response to a pure strategy averaged 45.67 (or \$0.91) per game. Of course, for diffuse priors these differences will diminish. For instance, the average payoff difference between the best and the next-best response to the uniform distribution was 11.12 (or \$0.22) per game, and the average payoff difference between the best and the worst response to the uniform distribution was 20.45 (or \$0.41) per game. We believe these differences are sufficient to invoke substantial effort. Further, it is implausible that any participant would have been able to infer the flat *ex post* payoff structure, since they received no feedback about anyone else's choices until after the experiment was completed; and if they did induce the payoff structure, then everyone should have adopted the easiest behavior (level-0), contrary to our findings (see also Merlo and Schotter, 1992).

Rather, the flat ex post payoff structure suggests that "evolutionary" forces may not be strong enough to drive out any of the types over a relatively short horizon. Even an "invading" perfectly precise RE type, who chooses the best-response on each game, while obviously doing better than all the boundedly rational types, does only slightly better relative to the statistical variation across games and other types.

## 6. Robustness Tests

How well will our estimated model predict out-of-sample? One way to address this important question is to reestimate the model on a subset of games, use these estimates to predict the behavior on the other games, and measure the stability of the parameter estimates, the goodness-of-fit, and the robustness of the posterior $\alpha$'s. The selection of an estimation subset (I) and a test subset (II) involves a number of tradeoffs. If the size of subset I is too small, then the decreased efficiency of the parameter estimates will obviously produce poorer test results. Further, if the composition of subsets I and II are very different (e.g., if subset I contains all of the dominance solvable games but none of the games with mixed-strategy NE), then the types might not be well-identified on subset I so it would be difficult to interpret the test results.

We decided that subset I should contain nine games to ensure reasonably efficient parameter estimates. Selecting the first (or last) three games for subset II was ruled out because there may be temporal effects that would confound the intended test. Instead, we selected games 4, 8, and 12, as an arbitrary but satisfactory temporal sampling. Further, this subset includes a variety of characteristics: game 4 has a strictly mixed-strategy NE, game 8 has a unique pure-strategy NE which is distinct from $b(P_0)$ and $b^2(P_0)$, and game 12 is dominance solvable.

We first estimated the model without RE types on subset I. We then evaluated the log-likelihood of subset I data using the parameter estimates shown in Table IV obtained with the full data set. Since the model has 11 parameters, the resulting likelihood-ratio statistic of 11.554 is distributed chi-square with 11 degrees of freedom, and hence has a $p$-value of 0.398. Thus we cannot reject the null hypothesis of no difference between the two sets of parameter estimates. We then estimated the parameters of the mixture model without RE types on subset II and compared the resulting log-likelihood value with that obtained by evaluating the log-likelihood of subset II data using the parameter estimates obtained from subset I. The resulting likelihood-ratio statistic of 8.690 again has the chi-square distribution with 11 degrees of freedom, and thus has a $p$-value of 0.650. Therefore, we conclude that the parameter estimates for the mixture model are quite robust to the selection of games.

In order to test how well the parameter estimates from subset I predict the choices observed in subset II, we computed the aggregated expected choice frequencies ($\hat{\pi}^{\mathrm{II}}$) for the subset II games for all 48 participants [Eq. (14)] using the parameter estimates of subset I. We then computed a goodness-of-fit statistic $\hat{\lambda}^{\mathrm{II}}$ using Eq. (15); since there are 3 games in subset II, the goodness-of-fit statistic has 6 degrees of freedom. For the entire sample of 48 participants for subset II, $\hat{\lambda}^{\mathrm{II}} = 7.969$, which is not significant at any commonly accepted level; hence, we conclude that the parameter estimates from subset I are able to predict the choices observed in subset II.

Next, we examined whether the data for the three games in subset II provide evidence which might reject the hypothesis that participants' types are the same across subsets I and II. Given the parameter estimates for subset I, we can compute the posterior Bayesian estimates of the type of each participant ($\alpha_l^{hI}$) as in Section 5.3. Let $k^{hI} \equiv \text{argmax}_l\{\alpha_l^{hI}\}$; i.e., $k^{hI}$ gives participant $h$'s most likely type based on the subset I data.

Now consider the following exercise. We compute the estimated choice probabilities for the subset II games, $P_{II}(\hat{\gamma}_l, \hat{\mu}_l, \hat{\varepsilon}_l)$, for each type $l = 1, \ldots, 4$, based on the subset I parameter estimates. Then we give these "type descriptions" to a neutral referee along with the subset II data and ask her whether participant $h$ is type $k^{hI}$. The first thing she would do is compute the likelihood of participant $h$'s three choices conditional on being each type, $P_{lII}^h(s_{II}^h)$, using the type descriptions and Eq. (7). Without further information, she would use a uniform prior and compute a posterior probability $\tilde{\alpha}_k^{hII} = P_{kII}^h(s_{II}^h)/\Sigma_l P_{lII}^h(s_{II}^h)$. If she is conservative, she might use 5% as the minimum for rejecting the null hypothesis; if more liberal, she would use a higher critical value, thereby rejecting in more cases.

To give our test some power, we choose 15% as our critical value and define the "15%-support set" $T_{15}^{hII}$ for each participant $h$ as the set of types for which the corresponding $\tilde{\alpha}_l^{hII} \geq 0.15$; i.e., $T_{15}^{hII} = \{l | \tilde{\alpha}_k^{hII} \geq 0.15, l = 0, \ldots, 4\}$. In other words, we do not reject the hypothesis that participant $h$ behaves as one type on both subsets of games if $k^{hI} \in T_{15}^{hII}$. Applying this test, we fail to reject for 35 participants; that is, over 72% of the participants appear to behave as one type.[18]

Note that our results do not mean that the other 13 participants actually switch between behavioral types. Eleven of these involve the worldly type (level-4), which we have already pointed out may mask several yet-to-be-identified types, in which case our mis-specification may manifest itself as apparent type switching. Another explanation for three of the participants involves dominated strategies. Choosing a dominated strategy in some game resulted in the participant (in all but one case) being identified as a level-0 type in Table V. However, with only five games with dominated strategies and only one or two instances of a dominated strategy being chosen (for a given participant) there is a good chance that both subsets I and II do not contain the choice of a dominated strategy, in which case the participant will be identified as a level-0 type in one subset but not in the other.

## 7. Conclusions

We have put forth a theory of boundedly rational strategic thinking in which human players are distinguished by their model of other players and their ability to identify optimal choices given their priors, yielding a

[18] We can increase the power of this test by asking whether participant $h$ is type $l$ for each type with $\alpha_l^{hl} \geq 0.25$. Doing so would yield only one additional rejection.

three-parameter family of probabilistic choice functions. Within this family we specified five archetypes. We designed and conducted an experiment to detect these archetypes and to estimate the parameters which define these archetypes, as well as a rational expectations type. The experimental evidence rejects the rational expectations hypothesis, but confirms the boundedly rational theory: i.e., based on statistical analysis of the experimental data, we conclude that the boundedly rational archetypes were definitely present in the population, but that no rational expectations types were present. Further, from posterior calculations, we were able to identify most participants' behavior across all 12 games as being observationally equivalent to one specific type. These results withstood a number of robustness tests based on subsets of games.

While our model appears to be the best currently available for describing the experimental data and predicting, we do not believe it is the final answer. The fact that the aggregate choices of the so-called worldly group were not well fit by the model predictions leads us to believe that the parameter estimates for this archetype mask considerable diversity in the population. Further research should look for additional archetypes that may be pooled in the worldly group of the current theory. However, we do know that the rational expectations type is not likely to be found. We also tested for (but did not formally report) the presence of a "perfect foresight" type,[19] and soundly rejected that hypothesis.

Our theory and empirical findings may serve as a complementary theory of initial conditions for dynamic learning theories.[20] Further research should investigate how players' models are updated after being given aggregate information about recent population choices.

## APPENDIX

### Participant Instructions

You are about to participate in an experiment about interdependent decision making. If you follow these instructions carefully and make good decisions you might earn a *considerable amount of money* which will be paid to you in *cash* at the end of the session.

The experiment will be conducted in two stages. In Stage I you and all other participants in this room will each make twelve decisions, and based on your combined choices, you will earn *tokens*. In Stage II, you will have the opportunity to receive *dollars* based on the

[19] A player has perfect foresight if his prior is equal to the actual frequency distribution of the *other* players in his experimental session (or in all three sessions). He then chooses a (perhaps imprecise) best response to this prior.

[20] For example, see Van Huyck *et al.* (1991, 1992), Crawford (1991, 1993), Boylan and El-Gamal (1993), and Roth and Erev (1993).

tokens you earned in Stage I. We describe Stage II first so you will understand how the tokens you earn affect the number of dollars you might win.

Description of Stage II

At the end of Stage I you will have earned between 0 and 100 tokens for each of twelve decisions. Each decision will be treated separately. The number of dollars you receive in Stage II will depend partly on the number of tokens you earned in Stage I and partly on chance.

Specifically, we have three ten-sided dice: one blue, one red, and one white. The blue die counts tens, the red die counts ones, and the white die counts tenths. For example, 5 blue, 7 red, and 3 white generate the number 57.3. The possible numbers are 00.0 to 99.9 in increments of 0.1. [Pause for demonstration.] During Stage II you will put the dice into a tumbler, shake and dump them into a cardboard tray, and record the number. If the number of tokens you earned for the first decision is GREATER THAN the dice-generated number, then you WIN $2.00. If the number of tokens you earned for the first decision is LESS THAN OR EQUAL TO the dice-generated number, you LOSE and get $0. You will repeat this throwing of dice until you have generated twelve numbers in all—one for each decision made in Stage I.

For example, if you earn 80 tokens for your seventh decision, and the seventh dice-generated number was 47, then you would win $2.00 for that decision. If you earn 35 tokens for your ninth decision, and the ninth dice-generated number was 66, then you would get $0 for that decision.

Observe that the number of tokens you earn in Stage I translates into the chance of winning $2.00 in Stage II for each decision. Thus, the more tokens you earn, the greater will be your chance of winning the $2.00 prizes.

Description of Stage I

During Stage I you and all other participants in this room will make twelve decisions. The tokens that you earn will depend on your choice and the choices of all other participants in this room.

Each decision that you face will be described by a *Decision Table* consisting of nine numbers arranged in three rows and three columns. Here is an example:

| | "T" | "M" | "B" |
|---|---|---|---|
| T | 30 | 20 | 70 |
| M | 40 | 80 | 0 |
| B | 60 | 100 | 50 |

The labels for the rows and columns will always be the same as in this example, but the numbers may differ from this example (but they will be the same for all participants). To indicate your decision, you will circle T, M, or B on the left-hand side of the Decision Table; one and only one choice is permitted. Depending on your choice and the choices of all the other participants, you will earn *tokens* as follows.

Let ($n_T$, $n_M$, $n_B$) denote the numbers of *other* participants in this session who chose T, M, and B respectively for this Decision Table. If you happen to choose T, then using the

T-row of the above Decision Table, you will score 30 points for each the $n_T$ players, 20 points for the $n_M$ players, and 70 points for each of the $n_B$ players. Your *token earnings* for this Decision Table will be your average point score: $(30n_T + 20n_M + 70n_B)/(n_T + n_M + n_B)$.

Similarly, if you had chosen M, then using the M-row your token earnings would have been $(40n_T + 80n_M + 0n_B)/(n_T + n_M + n_B)$; and if you had chosen B, your token earnings would have been $(60n_T + 100n_M + 50n_B)/(n_T + n_M + n_B)$.

We will now work through a numerical example of how your token earnings will be affected by your choices and the choices of all other participants. Consider the above Decision Table.

We will write down three columns of numbers with headings "T", "M" and "B" (see below). Next we write underneath the headings the numbers of other participants in this session who chose T, M, and B respectively for this Decision Table.

For example, suppose there were 50 *other* participants and 30 chose T, 10 chose M, and 10 chose B; i.e., $(n_T, n_M, n_B) = (30, 10, 10)$. Then write down "30, 10, 10" underneath the headings. [These round numbers make the arithmetic of our example easy, but for the real experiment we will use the actual choices of the participants in this room.]

Now suppose you chose T; then write down the T-row of the Decision Table underneath. Multiply each column, add the results, and divide by the total number of other participants (30 + 10 + 10) = 50:

| | | | | | |
|---|---|---|---|---|---|
| | | "T" | "M" | "B" | |
| Total of Others' Choices: | | 30 | 10 | 10 | (50) |
| Your Choice: | T | 30 | 20 | 70 | |
| | | 900 + | 200 + | 700 | = 1800; 1800/50 = ___36___. |

This means you would receive 36 tokens, giving you a 36% chance of winning $2.00 in Stage II.

*Practice Exercise 1.* Consider the same Decision Table. Suppose the total of other participants' choices are as before: $(n_T, n_M, n_B) = (30, 10, 10)$. How many tokens would you have earned if you had chosen M?

| | | | | | |
|---|---|---|---|---|---|
| | | "T" | "M" | "B" | |
| Total of Others' Choices: | | 30 | 10 | 10 | (50) |
| Your Choice: | M | 40 | 80 | 0 | |

__________.

*Practice Exercise 2.* Consider the same Decision Table. Suppose the total of other participants' choices are as before: $(n_T, n_M, n_B) = (30, 10, 10)$. How many tokens would you have earned if you had chosen B?

__________.

*Practice Exercise 3.* Given the same Decision Table and $(n_T, n_M, n_B) = (30, 10, 10)$, which choice would have earned you the most tokens:

T, M, or B ? (circle one)

**************************************************

With the next exercises, we will demonstrate how the choices of the other participants affect your token earnings. Consider the same Decision Table, but suppose that the other participants' choices were 10 T's, 0 M's, and 40 B's; i.e., $(n_T, n_M, n_B) = (10, 0, 40)$.

*Practice Exercise 4.* How many tokens would you have earned if you had chosen T?

| | | "T" | "M" | "B" | |
|---|---|---|---|---|---|
| Total of Others' Choices: | | 10 | 0 | 40 | (50) |
| Your Choice: | T | 30 | 20 | 70 | |

__________.

*Practice Exercise 5.* Consider the same Decision Table. Suppose the total of other participants' choices are as before: $(n_T, n_M, n_B) = (10, 0, 40)$. How many tokens would you have earned if you had chosen M?

| | | "T" | "M" | "B" | |
|---|---|---|---|---|---|
| Total of Others' Choices: | | 10 | 0 | 40 | (50) |
| Your Choice: | M | __ | __ | __ | |

__________.

*Practice Exercise 6.* Consider the same Decision Table. Suppose the total of other participants' choices are as before: $(n_T, n_M, n_B) = (10, 0, 40)$. How many tokens would you have earned if you had chosen B?

__________.

*Practice Exercise 7.* Given the same Decision Table and $(n_T, n_M, n_B) = (10, 0, 40)$, which choice would have earned you the most tokens:

T, M, or B ? (circle one)

******************************************************************************

CAUTION: The numbers used in these exercises were selected solely to make the arithmetic easy. They are *not* intended to suggest reasonable beliefs about how participants might respond if this Decision Table were used in the actual experiment.

Given *your beliefs* about the other participants' choices (i.e., $n_T, n_M, n_B$) in this session, you can compute your potential earnings for your choice via simple arithmetic (as above).

You know that the other participants are intelligent human beings like yourself, and that they are facing exactly the same Decision Table as you are.

**************************************************

Next I will give everyone a ten-minute Screening Test to ensure that everyone understands how his/her token earnings will depend on his/her choice and all other participants's choices in this session. The questions are like the practice exercises we just did. You must pass the screening test in order to participate in the experiment.

WAIT FOR MY SIGNAL BEFORE TURNING THE PAGE.

## SCREENING TEST

Consider the following Decision Table:

| | "T" | "M" | "B" |
|---|---|---|---|
| T | 40 | 0 | 60 |
| M | 50 | 10 | 100 |
| B | 20 | 30 | 50 |

*For all four questions, assume* the total of other participants' choices are: 20 T's, 20 M's, and 10 B's: i.e., $(n_T, n_M, n_B) = (20, 20, 10)$.

*Question 1.* What would be your token earnings if you had chosen T?

__________.

ASK ME TO CHECK YOUR ANSWERS BEFORE PROCEEDING.

*Question 2.* What would have been your token earnings if you had chosen M?

__________.

ASK ME TO CHECK YOUR ANSWER BEFORE PROCEEDING.

- - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -

YOU MUST ANSWER THE NEXT TWO QUESTIONS CORRECTLY TO PASS THIS TEST.

*Question 3.* What would have been your token earnings if you had chosen B?

__________.

*Question 4.* Which choice would have earned you the most tokens?

T M B (circle one)

ASK ME TO CHECK YOUR ANSWERS BEFORE PROCEEDING.

CONGRATULATIONS, you have passed the screening test and may now participate.

Stage I will take 36 minutes. Every participant in this room will be presented with twelve Decision Tables just like the examples we went through. There are four Decision Tables on each of the next three pages. DO NOT TURN THE PAGE until instructed to do so. Each Decision Table will also be displayed at the front of the room. You may use the blank space next to the Decision Tables as a scratch pad, and you may use a pocket calculator.

Paper clipped to the last page of this packet is a pink RECORD SHEET. Remove that sheet now. AFTER I give you the signal to start, you MUST DO two things for each Decision Table:

(1) Indicate your choice by *circling* T, M, or B on the left-hand side of each Decision Table. *One and only one choice is permitted.* You are not permitted to make more than one choice and you are not permitted to skip any Decision Table without making a choice. *Failure to make exactly one choice for each of the twelve Decision Tables will mean forfeiture of your token earnings for the entire experiment.*

(2) Record your choice in column 2 of the pink Record Sheet. Each Decision Table will be numbered 1 to 12. Be sure to record your choice in the proper place. This Record Sheet that will be used to determine your token earnings.

You will be told when 18, 10, 5, and 1 minutes remain in Stage I. You will also be issued a final 10 second warning. When the 36 minutes is up, you will be told to STOP, and pass in your Record Sheet.

Each Decision Table will be treated completely separately. That is, your token earnings for a specific Decision Table will depend on your choice and the choices of all other participants in this session for that Decision Table by *itself.* There is absolutely no linkage between different Decision Tables.

The following *Stage II instructions* will be repeated at the beginning of Stage II. I will need about 10 minutes to enter your choices into the computer. During this time, you will be asked to fill out the brief Post-Experiment Questionaire on the last page of this packet.

Also paper-clipped to the back of this packet is a blue EARNINGS SHEET on which you should record your choices for your personal record *after* turning in your pink Record Sheet. You will need this blue Earnings Sheet to convert token earnings into cash winnings.

After the choice data have been processed, the total responses for each Decision Table will be posted, and you will be invited to examine the Record Sheets to verify these totals.

You might also want to record these totals on your Earnings Sheet. [Then, you will have all the information necessary to verify the calculations of your token earnings.]

*One at a time*, each participant will be asked to come to the front station with the computer. You will be told privately and confidentially your token earnings for each Decision Table, and you will record these earnings on your Earnings Sheet. You will then proceed to the dice station, where you throw the dice and record the results on your Earnings Sheet.

Finally, you will proceed to the payment station, where your dollar winnings will be verified and paid to you in cash.

This is a serious scientific experiment, and as such, *no* TALKING, SIGHING, GROANING, LOOKING AROUND, or WALKING AROUND will be permitted. If you violate these rules, you will be asked to leave and all your potential winnings will be forfeited. If you have any questions during the experiment, raise your hand and I will come around. However, the scientific protocol will severely limit how I can respond after the experiment starts. Therefore, if you have any questions, you should ask them now.

WAIT FOR MY SIGNAL BEFORE TURNING THE PAGE.

## REFERENCES


AITKEN, M., ANDERSON, D., AND HINDE, J. (1981). "Statistical Modelling of Data on Teaching Styles," *J. Roy. Statist. Soc. A* **144**, 419–461.

ANSCOMBE, F., AND AUMANN, R. (1963). "A Definition of Subjective Probability," *Ann. Math. Statist.* **34**, 199–205.

BINMORE, K. (1987). "Modeling Rational Players, I and II," *Econ. Philosophy* **3**, 179–214; **4**, 9–55.

BOYLAN, R., AND EL-GAMAL, M. (1993). "Fictitious Play: A Statistical Study of Multiple Economic Experiments," *Games Econ. Behav.* **5**, 205–222.

BRANDTS, J., AND HOLT, C. (1992). "An Experimental Test of Equilibrium Dominance in Signalling Games," *Amer. Econ. Rev.* **82**, 1350–1365.

CAMERER, C., AND WEIGELT, K. (1988). "Experimental Tests of a Sequential Equilibrium Reputation Model," *Econometrica* **56**, 1–36.

COOPER, R., DEJONG, D., FORSYTHE, R., AND ROSS, T. (1990). "Selection Criteria in Coordination Games: Some Experimental Results," *Amer. Econ. Rev.* **80**, 218–234.

CRAWFORD, V. (1991). "An 'Evolutionary' Interpretation of Van Huyck, Battalio and Beil's Experimental Results on Coordination," *Games and Economic Behavior* **3**, 25–59.

CRAWFORD, V. (1993). "Adaptive Dynamics in Coordination Games," Department of Economics Discussion Paper 92-02R, University of California at San Diego.

EFRON, B. (1982). *The Jackknife, the Bootstrap, and Other Resampling Plans*. Philadelphia: Society for Industrial and Applied Mathematics.

EL-GAMAL, M., AND GRETHER, D. (1993). "Uncovering Behavioral Strategies: Likelihood-Based Experimental Data-Mining," Soc. Sci. Working Paper No. 850, California Institute of Technology.

EL-GAMAL, M. A., AND PALFREY, T. R. (1995). "Vertigo: Comparing Structural Models of Imperfect Behavior in Experimental Games," *Games and Economic Behavior* **8**, 322–348.

EL-GAMAL, M., MCKELVEY, R., AND PALFREY, T. (1993). "A Bayesian Sequential Experimental Study of Learning in Games," *J. Amer. Stat. Assoc.* **88**, 428–435.

El-Gamal, M., McKelvey, R., and Palfrey, T. (1994). "Learning in Experimental Games," *Econ. Theory* **4,** 901–922.

Everitt, B. S., and Hand, D. J. (1981). *Finite Mixture Distributions*. London: Chapman and Hall.

Friedman, D. (1993). "Equilibrium in Evolutionary Games: Some Experimental Results," Economics Department mimeo, University of California at Santa Cruz.

Hall, P. (1986). "On the Number of Bootstrap Simulations Required to Construct a Confidence Interval," *Ann. Statist.* **14,** 1453–1462.

Harrison, G. (1989). "Theory and Misbehavior of First-Price Auctions," *Amer. Econ. Rev.* **79,** 749–762.

Holt, D. (1993). "An Empirical Model of Strategic Choice with an Application to Coordination Games," Department of Economics mimeo, Queen's University, Canada.

McFadden, D. (1974). "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics* (P. Zaremba, Ed.) Academic Press.

McKelvey, R., and Palfrey, T. (1992). "An Experimental Study of the Centipede Game," *Econometrica* **60,** 803–836.

McKelvey, R., and Palfrey, T. (1994). "Quantal Response Equilibrium in Normal Form Games," mimeo, California Institute of Technology.

Merlo, A., and Schotter, A. (1992). "Theory and Misbehavior of First-Price Auctions: Comment," *Amer. Econ. Rev.* **82,** 1413–1425.

Nagel, R. (1993). "Experimental Results on Interactive Competitive Guessing," Discussion Paper 8-236, Sonderforschungsbereich 303, Universität Bonn.

Nelder, J. A., and Mead, R. (1965). "A Simplex Method for Function Minimization," *Computer J.* **7,** 308–313.

Ochs, J., and Roth, A. (1989). "An Experimental Study of Sequential Bargaining," *Amer. Econ. Rev.* **79,** 355–384.

Roth, A., and Erev, I. (1993). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," lecture presented at the Nobel Symposium on Game Theory.

Roth, A., and Malouf, M. (1979). "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Rev.* **86,** 574–1594.

Savage, L. (1954). *The Foundations of Statistics*. New York: Wiley.

Stahl, D. (1993). "Evolution of Smart$_n$ Players," *Games Econ. Behav.* **5,** 604–617.

Stahl, D., and Wilson, P. (1994). "Experimental Evidence of Players' Models of Other Players," *J. Econ. Behav. Organ.*, forthcoming.

Titterington, D. M., Smith, A. F. M., and Makov, U. E. (1985). *Statistical Analysis of Finite Mixture Distributions*. Chichester: Wiley.

Van Huyck, J., Battalio, R., and Beil, R. (1990). "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *Amer. Econ. Rev.* **80,** 234–248.

Van Huyck, J., Battalio, R., and Beil, R. (1991). "Strategic Uncertainty, Equilibrium Selection Principles, and Coordination Failure in Average Opinion Games," *Quart. J. Econ.* **106,** 885–910.

Van Huyck, J., Gillette, A., and Battalio, R. (1992). "Credible Assignments in Coordination Games," *Games Econ. Behav.* **4,** 606–626.