

Deductive versus inductive equilibrium selection: experimental results

Ernan Haruvy^{a,*}, Dale O. Stahl^b

^a School of Management, University of Texas at Dallas, Richardson, TX 75080, USA

^b Department of Economics, University of Texas at Austin, Austin, TX 78712, USA

Received 24 June 2002; received in revised form 10 September 2002; accepted 22 October 2002

Abstract

The debate in equilibrium selection appears to have culminated in the formation of two schools of thought: those who favor equilibrium selection based on rational coordination and those who favor zero-rationality adaptation. We examine four deductive selection principles and find that each fails to explain experimental data. We propose an inductive selection principle based on simple learning dynamics. Using out-of-sample maximum likelihood parameters, the predictive performance of one such dynamic is shown to be dramatically better than the deductive selection principles. However, this selection principle is not always definitive, since no dynamic is guaranteed to converge.

© 2003 Elsevier B.V. All rights reserved.

JEL classification: B41; C14; C51; C72; C90

Keywords: Deductive selection; Adaptive dynamics

1. Introduction

Equilibrium selection has been in the forefront of game theory in recent years, with the need for a salient selection method increasing as new economic and social problems involving multiple equilibria are being modeled. There are two main schools of thought in the area of equilibrium selection: On the one hand we have deductive selection—selection based on reasoning and coordination on focal points—and on the other hand we have inductive selection—selection based on adaptive dynamics. The debate between these two camps appears to have reached an impasse. Whereas existing deductive selection rules have been shown to do poorly in experiments (Van Huyck et al., 1990, 1991; Straub, 1995), inductive selection principles appear more promising.

* Corresponding author.

E-mail address: eharuvy@utdallas.edu (E. Haruvy).

The *deductive equilibrium selection* literature attempts to explain and predict which of the equilibria surviving refinements should be expected in different classes of games. A common conjecture is that decision-makers apply some deductive principle to identify a specific Nash equilibrium. One such deductive selection principle is *payoff-dominance* (Harsanyi and Selten, 1988, p. 81; Schelling, 1960, p. 291). Applying this principle, one expects the equilibrium outcome in a coordination game to be the highest Pareto-ranked equilibrium. The major limitation of payoff dominance (PD) lies in its failure to take into consideration off-equilibrium payoffs. To remedy this deficiency, equilibrium selection principles have been developed that are based on “riskiness,” the most famous of which is Harsanyi and Selten’s (1988) *risk dominance* (RD) selection principle.

Schelling was the first to note that the salience of a selection principle used in a particular game is largely an empirical question. His support of experimental methods came from his conviction that “some essential part of the study of mixed motive games is empirical.” Further, “the principles relevant to *successful* play, the *strategic* principles, the propositions of a *normative* theory, cannot be derived by purely theoretical means from a priori considerations” (Schelling, 1960, p. 162).

Experimental results [for prominent examples see Cooper et al. (1990); Van Huyck et al. (1990, 1991); henceforth, VHBB), Van Huyck et al. (1994, 1997); henceforth, VHCB), and Straub, 1995] do not appear to favor deductive principles. A possible explanation for the apparent failure of deductive principles is that they assume decision-makers possess beliefs consistent with some equilibrium without attempting to explain the process by which decision-makers acquire these equilibrium beliefs. Other experimental works [Stahl and Wilson (1994, 1995); henceforth, SW), Stahl (1996); Haruvy et al. (2001); henceforth, HSW), and Haruvy (2002)] reject the hypothesis that all experimental subjects generally begin with equilibrium beliefs. Hence, it would seem that an equilibrium outcome is generally not the result of choices made by decision-makers with equilibrium beliefs but rather the result of a dynamic process that begins with first period play by less-than-super-rational decision-makers.

Until recently, deductive selection principles that do not allow a role for the history of play or learning have dominated the equilibrium selection literature. The failure of deductive principles has shifted interest to learning and evolutionary dynamics as possible tools for equilibrium prediction. The basis for these *inductive selection principles* is the idea that in cases where decision-makers initially fail to coordinate on some equilibrium, repeated interaction may allow them to learn to coordinate. Having some experience in the game provides a decision-maker with observations that can be used to reason about the equilibrium selection problem in the continuation game. This experience may influence the outcome of the continuation game by focusing expectations on a specific equilibrium point.

Some experimental studies of games with multiple equilibria have found that relatively simple adaptive learning dynamics often yield good equilibrium predictions. In these experiments, knowledge of the initial distribution of play was sufficient to predict the equilibrium outcome (see VHBB, VHCB, and Roth and Erev, 1995). However, even with a good characterization of dynamics, one must specify the initial distribution of play before predicting the final outcome. Recent research (Haruvy and Stahl, 2000) has attempted to fill this gap by studying alternative theories (a priori specifications) of initial conditions. They find that specifying uniform initial conditions for “period 0” (i.e. a fictitious period prior to the actual

first period of play) and using the dynamic model to predict play for period 1 onward is a robust and parsimonious specification that fits the dynamics quite well.

Stahl (1999) conducted a horse race among seven action-reinforcement learning models and found that a simple four-parameter *logit best-reply with inertia and adaptive expectations* (LBRIAE) dynamic outperformed all others both in sample and out-of-sample by several measures. We therefore focus on the LBRIAE dynamic model in this paper as a candidate for an inductive equilibrium selection principle; if the LBRIAE dynamic converges, then we call the limit point the *LBRIAE equilibrium*.

The Harsanyi and Selten tracing procedure has both deductive and inductive features. Their algorithm adjusts arbitrary prior beliefs into equilibrium beliefs through gradual movement in the direction of best response to the prior beliefs. It is important to recognize, however, that their underlying dynamic process occurs in the mind of the player before the first period of play, so it is independent of empirical histories in a given game. Further, unlike LBRIAE dynamics, dominated strategies have no effect on their predictions—an implication that has been strongly refuted by experimental data (e.g. Cooper et al., 1990). Nonetheless, we adopt the spirit of their approach, suggesting simple initial conditions and moving in a dynamic manner, to arrive at an ex-ante prediction based on game properties alone.

We describe four deductive selection principles in Section 2 and our inductive approach in Section 3. Section 4 describes some simple games that test existing notions of deductive selection against our proposed alternative and the experimental procedure. Section 5 describes the results, and Section 6 concludes.

2. Deductive equilibrium selection principles

In this section we briefly review the main deductive selection principles in the literature: payoff dominance, security (SEC) and risk dominance. The premise behind the deductive selection principles is that players choose an action from the set of Nash equilibrium actions according to various criteria. If all players apply the same criterion, the equilibrium outcome can be predicted without any consideration of dynamics.

2.1. Payoff dominance

The payoff dominance principle relies on the idea that “rational individuals will cooperate in pursuing their common interests if the conditions permit them to do so” (Harsanyi and Selten, 1988, p. 356). In the symmetric normal-form games we study, the payoff dominant equilibrium corresponds to the Nash equilibrium action with the largest diagonal payoff. Experimental studies by Cooper et al. (1990, 1992), VHBB (1990, 1991) and Straub (1995) on coordination games provide substantial evidence that players often fail to coordinate their actions to obtain a Pareto-optimal equilibrium in experimental settings.

2.2. Security

A *secure action* is one that maximizes the minimum possible payoff (Van Huyck et al., 1990). Thus, when each act is appraised by looking at the worst state for that act, the secure

action is the action with the best worst state. This idea is the pure-strategy version of [Von Neumann and Morgenstern's \(1947\)](#) maximin criterion. It is important to note that in games with non-Nash actions, there is no reason to assume that the secure action should be in the support of some Nash equilibrium. Therefore, to make the security criterion an equilibrium selection principle, it must be modified to exclude actions that are not in the support of some equilibrium. We restrict the security criterion to equilibrium actions by defining the *secure equilibrium action* as that equilibrium action that satisfies

$$\arg \max_{k \in \text{NE}} \min_{j \in \text{NE}} U_{kj} \quad (1)$$

where U is a $J \times J$ matrix of game payoffs for the row player in a given game and NE denotes the set of Nash equilibrium actions. This specification applies the security criterion to the game after the deletion of non-equilibrium actions. In accordance with this restriction, the *security* selection principle is an equilibrium selection principle that predicts the maximin action after restricting attention to the set of equilibrium actions.

2.3. Risk dominance

The Harsanyi and Selten *risk dominance* selection criterion is concerned with pair-wise comparisons of Nash equilibria. The equilibrium with the highest *Nash-product* is selected out of each pair, where the term Nash-product refers to the product of the deviation losses of both players at a particular equilibrium. Unfortunately, there are difficulties when attempting to apply this definition to general $n \times n$ games with more than two equilibria because the pair-wise risk dominance (PRD) relation is not necessarily transitive.

One solution is to redefine risk dominance in accordance with Harsanyi and Selten's *heuristic justification*, in which selection of an equilibrium results from postulating an initial state of uncertainty where the players have uniformly distributed second order beliefs (i.e., each player believes that the other players' beliefs are uniformly distributed on the relevant space of priors). Briefly, given a symmetric $n \times n$ game with payoff matrix U , let NE denote the set of Nash equilibrium actions, and let Δ^{NE} denote the simplex on NE. For each $j \in \text{NE}$, define q_j^{RD} as the relative proportion of Δ^{NE} for which action j is the best response to some belief in Δ^{NE} . Then the action $k \in \text{NE}$ that maximizes $U_k q^{\text{RD}}$ (where U_k is the k th row of the payoff matrix) is the risk-dominant NE action. This solution coincides with the pair-wise definition in 2×2 games and ensures transitivity of the risk dominance relation in symmetric $n \times n$ games. We shall refer to this extension simply as *risk dominance*.

To illustrate the difference between the two concepts, consider game 1 in [Fig. 1](#). There are three equilibria in this game. The Nash product for equilibrium A in pair (A, B) is 10^2 and for equilibrium B in that pair it is 20^2 . Hence B dominates A. The Nash product for equilibrium A in pair (A, C) is 30^2 and for equilibrium C in that pair it is 10^2 . Hence A dominates C. The Nash product for equilibrium B in pair (B, C) is 60^2 and for equilibrium C in that pair it is 50^2 . Hence B dominates C. Therefore, B is the pair-wise risk dominance solution.

For the RD solution, the RD prior is $q^{\text{RD}} = (0.73, 0.22, 0.05)$, corresponding to the relative areas of the best response regions for each equilibrium, and the expected payoffs

Game 1 (5 ses)

11806

A	70	60	90
B	60	80	50
C	40	20	100

L1, SEC, RD,
LBRIAE

PRD

PD

Game 4 (5 ses)

11182

A	70	30	20
B	60	60	30
C	45	45	40

PD

L1, RD,
LBRIAE

SEC

Game 13 (7 ses)

2211142

A	60	60	30
B	30	70	20
C	70	25	35

L1

PD, RD

SEC,
LBRIAE

Game 14 (5 ses)

38832

A	50	0	0
B	70	35	35
C	0	25	55

DOM

L1, SEC,
LBRIAE

PD, RD

Game 19 (5 ses)

111085

A	80	60	50
B	60	70	90
C	0	0	100

PRD

L1, SEC, RD
LBRIAE

PD

Key: PD = Payoff dominant Nash equilibrium strategy
RD = Risk dominant Nash equilibrium strategy
PRD = Pair-wise risk dominant strategy (only indicated when distinct from RD)
SEC = Security Nash equilibrium strategy
L1 = Level-1 Strategy
LBRIAE = large-population limit distribution with no trembles
Underlined numbers are the aggregate choices.

Fig. 1. The game matrices, number of sessions and aggregate choices.

are 68.8, 63.9 and 38.6 for A, B, and C, respectively. Therefore, A is the best response to the RD prior and is the prediction for the RD principle.

In the games we study, the Harsanyi–Selten pair-wise definition of risk dominance yields a unique solution that we denote as *pair-wise risk dominance*. Harsanyi and Selten also introduce a tracing procedure as a risk dominance approach for more general games. In our games, the Harsanyi–Selten tracing procedure would pick the PRD equilibrium. It is important to note that the selection principle promoted by Harsanyi and Selten would in fact select the unique PD equilibrium over any of the risk dominance concepts. We nonetheless isolate the risk dominance notion as a principle worth investigating for its own merits.

3. Inductive selection principles

By inductive selection principles we refer to dynamic models and their limit points. We begin by addressing several methodological issues concerning the predictions of a dynamic model. We then present the LBRIAE model as an example and our preferred inductive principle.

3.1. Inductive processes as selection principles

Though several models of dynamics have been proposed in the literature in the context of coordination in games with multiple equilibria, few authors have focused on dynamic models as a solution to the equilibrium selection problem. One exception is [Van Huyck et al. \(1997\)](#) who studied adaptive behavior in a generic game with multiple Pareto ranked equilibria. They found that (i) behavior diverged at the separatrix—the border separating the basins of attraction for each equilibrium—predicted by the fictitious play dynamic, and (ii) the equilibrium selected was sensitive to small differences in initial conditions. However, they made no characterization of the appropriate initial conditions. Obviously, such sensitivity is an impediment to using inductive processes to define selection principles.

Unlike the standard dynamic literature reliance on one-period-ahead measures of likelihood, the focus in the application of dynamics to equilibrium selection is on *T-period-ahead* prediction, or the ex-ante prediction prior to the start of the game of the period *T* frequency of choice. To compute the exact theoretical probability distribution for such a *T*-period-ahead prediction, we would need to integrate out the $T - 1$ periods prior to *T*. Although such a feat may be impossible analytically, it can be approximated to any desired degree of accuracy by simulating a large number of paths of play.

Given a finite population and a positive probability (ϵ) of trembles, every *T*-period path has a strictly positive probability. Therefore, our integration procedure will put positive probability on every stable Nash equilibrium. This indeterminacy is clearly a drawback to an inductive selection principle. One way to generate more definitive predictions is to simulate paths of play for a large population (thereby reducing the multinomial variance), and to reduce the probability of trembles, taking the limit as the population size increases indefinitely and ϵ goes to zero. The limit dynamics are those of a deterministic first-order difference equation on the simplex. Starting with initial conditions for the choice frequency $p(0)$, if the limit dynamics converge, then the limit point is the inductive selection principle's prediction.

There are two caveats of this limit approach. *First*, the limit predictions are not necessarily reliable for finite populations and positive trembles. For small populations with non-negligible trembles, historical accidents could have a permanent effect on the long-run outcome by bumping the path out of one basin of convergence into another. Hence, when attempting predictions for small populations, it would be safer to use simulations for the actual population size and tremble likelihood.

Second, the predictions may be highly sensitive to the initial conditions $p(0)$. To see this, consider a game for which $p(0)$ is very close to a separatrix. Slightly perturbing the initial conditions so they lie on the other side of the separatrix will result in dramatically different final outcomes for the limit dynamics. Thus, unless one is extremely confident

in the specification of initial conditions, one should be concerned about the robustness of the limit results to initial conditions. While using simulations for the actual population size and empirically measured tremble probability will mitigate this problem somewhat, a better approach would be direct sensitivity analysis: e.g., draw initial priors from a multinomial distribution as if period 0 were real.

3.2. The LBRIAE inductive selection principle

Stahl (1999) proposed the *logit best-reply with inertia and adaptive expectations* model. The population is assumed to be comprised of two types. One type either sticks with last period's choice or imitates the most recent empirical frequency of the whole population, $p(t-1)$. Given a fixed propensity to stick or imitate, the resulting behavior of this type is a first-order dynamic process that has the same structural form as an adaptive expectations process for beliefs. The second type is assumed (i) to have beliefs given by the adaptive process of the first type (as if they believe everyone else is of the first type), and (ii) to choose a noisy (logistic) best-reply to this belief. To accommodate trembles by all types, the probability choice function is mixed with the uniform distribution over the actions. Furthermore, unlike the standard assumption of uniform initial conditions for period 1, LBRIAE imposes the uniformity assumption on a fictitious period 0 and uses the dynamic model to predict first period behavior.

Defining the LBRIAE prediction as the limit of the large population dynamics as the tremble probability goes to zero, the prediction is a logit-response equilibrium of the game (McKelvey and Palfrey, 1995) that will depend on the predetermined values of the LBRIAE parameters. We hasten to point out that there is no guarantee that the limit dynamics will converge.

We agree with VHCB that simple “better-response” dynamics should be expected to predict well for many games with multiple equilibria, and we deem the four-parameter LBRIAE model to belong to this class. Moreover, it appears that the tremble structure and the herd behavior of this model result in a much better fit of experimental data than other leading models (Stahl, 1999). While the final equilibrium outcome for most of our games is predicted equally well by all leading dynamic models, we will see in Section 5 that LBRIAE outperforms Anderson et al. (2001); Roth and Erev (1995), and Camerer and Ho (1999) for one of the games. Hence, we focus on the inductive selection principle derived from the LBRIAE model. We nevertheless briefly describe the other three learning models. The reader should refer to the original paper for a more detailed description.

3.3. Other learning models

LBRIAE could be classified as a belief learning model since players are modeled as forming beliefs about the probabilities of opponent choices based on information about past opponent choices, irrespective of the choice actually played by the individual player forming the beliefs. A simpler version of belief learning, which we call the logit form of replicator dynamics, is derived from the idea of stochastic fictitious play with inertia (see Fudenberg & Levine, 1998; Anderson et al., 2001). The choice rule is stochastic, and beliefs are updated with depreciation of past beliefs by a constant.

In contrast to belief learning, Roth and Erev presented what has come to be known as the reinforcement learning model. In their model, players update propensities based only on realized payoffs to the strategies actually played. Camerer and Ho's model is presented as a hybrid approach between belief and reinforcement learning, where the two schools are described as differing on the weight given by human players to foregone payoffs in updating propensities. As noted by Camerer and Ho, belief learning does not distinguish foregone from realized payoffs in the updating of propensities, whereas players in reinforcement learning models do not pay any attention to foregone payoffs. The idea behind the EWA model is that players evaluate the performance of each possible action in the last period and update their propensities to use each action accordingly. However, the action actually played by each player receives greater attention in the evaluation process. Hence, actions are reinforced according to past performance, but actions actually played receive some additional reinforcement.

4. The games and experimental procedure

We selected five games that discriminate among the deductive equilibrium selection principles of payoff dominance, risk dominance, pair-wise risk dominance, and security. For the LBRIAE selection principle and all other learning models, we use parameters estimated in Stahl (1999) for a totally different data set and produce 10,000 simulations for a large population and vanishing trembles.¹ The five games are (using HS99 numbering): 1, 4, 13, 14 and 19, shown in Fig. 1. These are symmetric normal-form games, so all players read the payoff matrix as “row” players. That is, the choices A, B, and C are row choices, where A is the top row, B is the middle row, and C is the bottom row. The numbers in each row give the payoff if everyone else chooses A, B or C, respectively. These five games all have the property that each of the selection principles makes a unique prediction as indicated in Fig. 1.² The aggregated choices are displayed as underlined numbers above each payoff matrix in Fig. 1.

A “mean-matching” protocol was used.³ In each period, a participant's token payoff was determined by her choice and the percentage distribution of the choices of all other participants, $p(t)$, as follows. The row of the payoff matrix corresponding to the participant's choice was multiplied by the vector of choice distribution of the other participants. Token payoffs were in probability units for a fixed prize of US\$ 2.00 per period of play. In other

¹ Since the equal-probable point is not close to any separatrix for these games, the limit predictions are robust to the initial conditions.

² While the limit point of the LBRIAE model is a logit equilibrium, since the estimated precision of the logit best-replies is high enough to put the limit point very close to a pure-strategy Nash equilibrium, for the purpose of comparisons with the deductive selection principles, we identify the LBRIAE prediction as that closest pure-strategy Nash equilibrium.

³ While pure theory would hold that playing against a single opponent randomly selected from the population is equivalent to playing against the field, the latter protocol makes asymmetric equilibria highly unlikely. Friedman (1996) 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. Recently, the mean-matching protocol has been found to promote strategic and individualistic (money-maximizing) behavior (Stahl and Haruvy, 2002a, 2002b).

words, the token payoff for each period gave the percentage chance of winning US\$ 2 for that period. The lotteries that determined final monetary payoffs were conducted following the completion of both runs using dice. Specifically, a random number uniformly distributed on $[00.0, 99.9]$ was generated by the throw of three 10-sided dice. A player won US\$ 2.00 if and only if his token payoff exceeded his generated dice number. Payment was made in cash immediately following each session.

Participants were seated at private computer terminals separated so that no participant could observe the choices of other participants. The relevant game, or decision matrix, was presented on the computer screen. Each participant could make a choice by clicking the mouse button on any row of the matrix, which then became highlighted. In addition, each participant could make hypotheses about the choices of the other players. An on-screen calculator would then calculate and display the hypothetical payoffs to each available action given each hypothesis. Participants were allowed to make as many hypothetical calculations and choice revisions as time permitted. Following each time period, each participant was shown the aggregate choices of all other participants and could view a record screen with the history of the aggregate choices of other participants for the entire run.

5. Results

We first examine the aggregate final-period choices and compute the proportion of those choices (aggregated over all experimental sessions of a game) that are consistent with equilibrium selection principle P , where $P \in \{PD, RD, PRD, SEC, LBRIAE\}$. Because we have a different number of experimental sessions for the various games, we first average results for each game over the sessions of that game and finally take the simple average of these averages (Table 1).

We observe that PD performs worse by this criterion, since only 8.4% of the aggregate final-period choices are consistent with the PD principle. The LBRIAE selection principle⁴ clearly performs best by this criterion, since at least 70% of the aggregate final-period choices are consistent with the LBRIAE principle. While RD and SEC perform well above the 50% level, there are games for which these principles perform dismally (13 and 4, respectively). To see the robustness of these results across the games, note that the LBRIAE column weakly dominates the other four columns and that the RD column weakly dominates PR and PRD; thus, these rankings are invariant to any distribution across these games.

An alternative criterion for evaluating selection principles is by “outcomes” determined on a session-by-session basis. We say that *the final outcome is x* in session i of a game when at least 75% of the final-period choices are x ; here x stands for the action corresponding to

⁴ It is important to stress again that LBRIAE is used here with predetermined values from another set of games in another paper (Stahl, 1999), as are all other learning models considered here. As such, LBRIAE and the other learning models are selection principles in the purest definition of the term, rather than fitted on the data. The latter approach would put the inductive principles at an unfair advantage over the deductive principles (which cannot be fitted).

Table 1

Proportion of final-period choices consistent with equilibrium selection principle

	PD	RD	PRD	SEC	LBRIAE
Game 1	0.048	0.952	0	0.952	0.952
Game 4	0.008	0.975	0.975	0.017	0.975
Game 13	0.063	0.063	0.063	0.811	0.811
Game 14	0.260	0.260	0.260	0.715	0.715
Game 19	0.040	0.871	0.089	0.871	0.871
Average	0.084	0.624	0.277	0.673	0.865

Table 2

Proportion of final outcomes consistent with equilibrium selection principle

	PD	RD	PRD	SEC	LBRIAE
Game 1	0	1	0	1	1
Game 4	0	1	1	0	1
Game 13	0	0	0	0.857	0.857
Game 14	0.2	0.2	0.2	0.8	0.8
Game 19	0	1	0	1	1
Average	0.04	0.64	0.24	0.731	0.931

a particular equilibrium selection principle for that game. We then compute the proportion of the experimental sessions for which x was the final outcome (Table 2).

Once again we see that PD performs worst (4%) and LBRIAE best (93%). All the selection principles except LBRIAE completely miss the final outcome for at least one game in our data, while the lowest performance of LBRIAE is 80%. The performance of SEC and LBRIAE differ for only game 4, where action C is the unique SEC action, but B is the final outcome for all sessions run, and B is the RD and LBRIAE solution. RD however predicts none of the game 13 outcomes and only 20% of game 14 outcomes. Again, the LBRIAE column weakly dominates the other four columns, and RD weakly dominates PD and PRD, so these rankings are invariant to the distribution over games.

We see that both performance criteria rank the principles the same. LBRIAE, as a representative learning model, is a clear winner over the deductive principles. It missed only two outcomes: one for game 13, which all selection principles missed because the path did not converge, and one for game 14 that crossed the separatrix.

While most of the alternative learning models would make the same predictions as LBRIAE, game 19 presents a discriminating test. All five sessions of game 19 converged to the B outcome, which is predicted by the LBRIAE model. However, using Stahl's (1999) ML parameter estimates, Anderson, Goeree and Holt's logit form of replicator dynamics⁵ predicts 70% As, Roth and Erev's reinforcement learning predicts 52% As, and Camerer and Ho's EWA predicts 14% As. Ultimately, however, the purpose of this paper is not to present a horse race between learning models. The superiority of LBRIAE over competing

⁵ Anderson et al. (2001) discuss equilibrium selection implications.

learning models lies mainly in one game out of five, and this could hardly be considered sufficient evidence to rule one learning model better than others. The interested reader should see Stahl (1999) for a comprehensive horse race between learning models that includes LBRIAE as one of the candidates.

6. Conclusions

Our results confirm studies like VHBB and Straub in indicating that none of the mainstream deductive principles give reliable predictions. Though consistent with the spirit of the tracing procedure of Harsanyi and Selten, LBRIAE does not ignore dominated actions and may therefore produce very different suggestions from the tracing procedure. However, we hasten to point out that the LBRIAE process may not converge for some games. On the other hand, it is not clear that human play will converge in games where the LBRIAE process does not converge, and since the LBRIAE model fits the empirical dynamics quite well, the stochastic prediction implicit in the model may be about as good as we can hope for.⁶

When predicting final outcomes, we have assumed that there is no previous experience with a particular game (or saliently similar game) among the experiment participants. Therefore, the LBRIAE model invokes the principle of insufficient reason to specify a uniform prior for a fictitious period before the first period of the experiment. Furthermore, if past experience came from the same informational/institutional environment as the current experiment, then the farther into the past we go, the more reasonable is the assumption of a uniform prior sometime in the past, and hence the more reasonable is the LBRIAE prediction using a uniform prior.

On the other hand, finite-population stochastic dynamics can result in substantially different short-run outcomes (e.g. our Game 14). If the *same population* were to continue to play this game, the principle of insufficient reason would no longer apply, but instead it would be natural to specify the initial conditions as a function of the common experience (perhaps the historical empirical frequency distribution mixed with the uniform distribution). If the history-dependent initial conditions were to fall in a different basin of attraction, then the LBRIAE prediction would differ accordingly. Thus, the LBRIAE model can be easily adapted to accommodate relevant common experience.

In typical university-based experiments with thoroughly mixed (as opposed to isolated) subject populations and with abstract games that are not obviously similar to naturally occurring games for which there is substantial common experience, the LBRIAE model with a uniform prior (and perhaps some diversity) is likely to predict very well. However, in cross-cultural studies (e.g., Roth et al., 1991), in carefully designed experiments with experienced subjects, and in experiments with “natural” context-rich games, the predictive performance of the LBRIAE model may be improved by incorporating the relevant common experience into the specification of initial conditions.

⁶ The focus here is on simple dynamics. We believe that more sophisticated dynamics, such as the rule learning model of Stahl (2000), can do better.

Appendix A. The formal LBRIAE model

Let $q(t, \theta)$ denote the expected probability of play in period t based on the history of play up to and including period $t - 1$, and let $p(t - 1)$ denote the actual frequency of play in period $t - 1$. Then, the one-parameter adaptive expectation model specifies that

$$q(t, \theta) = \theta q(t - 1, \theta) + (1 - \theta) p(t - 1), \quad (2)$$

where, in accordance with the principle of insufficient reason, $q(0, \theta)$ and $p(0)$ are specified as the uniform distribution over the actions.⁷

It is assumed that a proportion δ of the population behaves according to Eq. (2), which can be interpreted as “herd” behavior in which with probability θ the past action will be repeated and with probability $1 - \theta$ the recent past will be mimicked. The proportion $1 - \delta$ of the population chooses a logit best-reply to Eq. (2). That is, letting $b(q, \nu)$ denote the logit best-reply to q with precision ν , then the probability choice function for period t conditional of history h^t is given by

$$\varphi(t|h^t) = \delta q(t, \theta) + (1 - \delta) b(q(t, \theta), \nu). \quad (3)$$

To accommodate trembles this probability choice function is mixed with the uniform distribution over the actions (denoted p^0):

$$\varphi^*(t|h^t) = (1 - \varepsilon) \varphi(t|h^t) + \varepsilon p^0. \quad (4)$$

Thus, this LBRIAE model has four parameters: ν , δ , θ , and ε . The Stahl (1999) maximum-likelihood parameter estimates are (0.3955, 0.3258, 0.2507, 0.0530), respectively.

Then, the probability of observing empirical frequency $p(t)$ in period t in a population of size n is a multinomial distribution, call it $\Phi^n(p(t)|h^t)$, with $\varphi^*(t|h^t)$ as the underlying probabilities.

To generate a T -period-ahead prediction of the outcome for the T th period, we need to integrate out the $T - 1$ periods prior to T :

$$\begin{aligned} \Phi^n(p(T)) &= \int \cdots \int \Phi^n(T|p(T - 1), \dots, p(1)) \\ &\quad \times \cdots \times \Phi^n(1|p(0)) dp(1) \cdots dp(T - 1). \end{aligned} \quad (5)$$

Note that while the domain for φ^* is the discrete set of the J actions in the game, the domain for Φ^n is the J -dimensional simplex.

References

- Anderson, S., Goeree, J., Holt, C., 2001. Minimum-effort coordination games: stochastic potential and logit equilibrium. *Games and Economic Behavior* 34, 177–199.
- Camerer, C., Ho, T., 1999. Experience-weighted attraction learning in normal form games. *Econometrica* 67, 827–874.
- Cooper, R., DeJong, D., Forsythe, R., Ross, T., 1990. Selection criteria in coordination games: some experimental results. *American Economic Review* 80, 218–233.

⁷ If there were some prior history for the game, then there might be a reason for a non-uniform prior.

- Cooper, R., DeJong, D., Forsythe, R., Ross, T., 1992. Communication in coordination games. *Quarterly Journal of Economics* 107, 739–771.
- Friedman, D., 1996. Equilibrium in evolutionary games: some experimental results. *The Economic Journal* 106, 1–25.
- Fudenberg, D., Levine, D., 1998. *The Theory of Learning in Games*. MIT Press, Cambridge.
- Harsanyi, J., Selten, R., 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge.
- Haruvy, E., 2002. Testing modes in beliefs. *Journal of Mathematical Psychology* 46, 88–109.
- Haruvy, E., Stahl, D., 2000. Robust initial conditions for learning dynamics. Working paper, University of Texas at Austin.
- Haruvy, E., Stahl, D., Wilson, P., 2001. Modeling and testing for heterogeneity in observed strategic behavior. *Review of Economics and Statistics* 83, 146–157.
- McKelvey, R., Palfrey, T., 1995. Quantal response equilibria for normal form games. *Games and Economic Behavior* 10, 6–38.
- Roth, A., Erev, I., 1995. Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8, 164–212.
- Roth, A., Prasnikar, V., Okuno-Fujiwara, M., Zamir, S., 1991. Bargaining behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: an experimental study. *American Economic Review* 81, 1068–1095.
- Schelling, T.C., 1960. *The Strategy of Conflict*. Harvard University Press, Cambridge.
- Stahl, D., 1996. Boundedly rational rule learning in a guessing game. *Games and Economic Behavior* 16, 303–330.
- Stahl D., 1999. A horse race among action-reinforcement learning models. Working paper, University of Texas at Austin.
- Stahl, D., 2000. Rule learning in symmetric normal-form games: theory and evidence. *Games and Economic Behavior* 32, 105–138.
- Stahl, D., Haruvy, E., 2002a. Level-n bounded rationality on a level playing field of sequential games. Working paper, University of Texas at Austin.
- Stahl, D., Haruvy, E., 2002b. Other-regarding preferences: Egalitarian warm glow, empathy, and group size. Working paper, University of Texas at Austin.
- Stahl, D., Wilson, P., 1994. Experimental evidence of players' models of other players. *Journal of Economic Behavior and Organization* 25, 309–327.
- Stahl, D., Wilson, P., 1995. On players' models of other players: theory and experimental evidence. *Games and Economic Behavior* 10, 213–254.
- Straub, P.G., 1995. Risk dominance and coordination failures in static games. *The Quarterly Review of Economics and Finance* 35, 339–363.
- Van Huyck, J., Battalio, R., Beil, R., 1990. Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review* 80, 234–248.
- Van Huyck, J., Battalio, R., Beil, R., 1991. Strategic uncertainty, equilibrium selection principles, and coordination failures in average opinion games. *Quarterly Journal of Economics* 106, 885–911.
- Van Huyck, J., Cook, J., Battalio, R., 1994. Selection dynamics, asymptotic stability and adaptive behavior. *Journal of Political Economy* 102, 975–1005.
- Van Huyck, J., Cook, J., Battalio, R., 1997. Adaptive behavior and coordination failure. *Journal of Economic Behavior and Organization* 32, 483–503.
- Von Neumann, J., Morgenstern, O., 1947. *Theory of Games and Economic Behavior*. Princeton, New Jersey.