Robust Initial Conditions for Learning Dynamics

A preliminary version

Ernan Haruvy
Faculty of Industrial Engineering and Management
Technion-- Israel Institute of Technology

Dale O. Stahl
Department of Economics
University of Texas

August, 2000

Abstract

Although a *complete* dynamic model includes the equations of motion as well as the initial conditions, many of the learning dynamics in the literature do not entail a serious theory of initial conditions. This paper explores and evaluates five approaches to complete learning models with a theory of initial conditions. We find that (1) the initial period's fit is sensitive to the model of initial conditions chosen, with some initial condition models consistently superior to others over games. (2) Looking at periods 2 onward, leading dynamic models are robust to initial condition models according to one-period-ahead goodness-of-fit criteria. (3) Longer horizon measures of fit provide the insight that models of initials conditions with initially high diversity sometimes outperform better specified models of initial conditions
.

Keywords: Adaptive Dynamics; Initial Conditions.

1. INTRODUCTION

The common element to theories of adaptive dynamics is that given initial conditions, the population will move in the direction of the (possibly noisy) best response. With the exception of this common element, different theories vary along numerous dimensions, such as attention to foregone payoffs, appropriate discounting of beliefs or propensities, mapping between beliefs or propensities and choice, level of sophistication, foresight, transference between games, heterogeneity, and various other aspects. Several works have attempted to conduct model comparisons (Camerer, 1999; Feltovich, 2000; Stahl, 1999; Erev and Haruvy, 2000) and highlight some of the differences between models. Attention in these comparisons has focused on various differences between models, but one aspect which has been all but ignored in learning models is the treatment of *initial conditions*. Some leading works on dynamics circumvent initial conditions by plugging in the actual empirical frequency in the first period (e.g., Camerer and Ho, 2000; Van Huyck, Cook, and Battalio, 1997). At the other extreme, one could apply the principle of insufficient reason to justify the uniform distribution as the default initial condition for the first period. While this approach will yield poor predictions for the first period, it could yield robust T-period ahead predictions (as claimed by Roth and Erev, 1998).

In parallel, works on one-shot games (Stahl and Wilson, 1994, 1995, Haruvy and Stahl, 1999, and Costa-Gomes et al., 1998), while rigorously characterizing the choice distribution, merely allude to their extension into dynamic predictions. Haruvy and Stahl (1999; hereafter HS99) tested a large class of models of initial play, including the major equilibrium selection principles. They concluded that the best predictor of initial play was given by a heterogeneous model with boundedly rational players. Using the parameters estimated by HS99, forecasts of first-period play may be generated for any symmetric normal-form game and used as the initial conditions for dynamics. We find that indeed the HS99 predictions for the (out-of-sample) data considered herein are quite accurate and the best predictions among the alternatives considered. Unfortunately, the HS99 model is parameter-intensive and hence cumbersome to re-estimate on new data sets. In HS99, the best homogeneous, one-parameter model of initial play was the simple logistic "level-1" model. Because of its simplicity and reliability, we also consider this level-1 model herein. In addition, we investigate an approach proposed by Stahl (1999) that adds a fictitious period 0 in which players play uniformly. We name this approach “uniform period 0.” To distinguish this from uniform initials in period 1 (such as those in Erev-Roth, 1995; Roth-Erev, 1998), we call the latter approach "uniform period 1."

Our study entails two conceptual dimensions: (1) initial conditions, and (2) dynamic models. We note that different initial condition models are ranked differently depending on both the dynamic and the measure used. We discuss and seek explanations for such discrepancies.

2. FIVE MODELS OF INITIAL CONDITIONS

Five approaches to dealing with initial conditions arise from extant works: (1) plugging in actual initials (subsection 2.1), (2) uniform period 1 initials (subsection 2.2), (3) mixtures of heterogeneous types (subsection 2.3), (4) level-1 (subsection 2.4), and (5) uniform period 0 (subsection 2.5). We briefly describe each.

2.1. Actual Empirical Frequency as Initial Condition.

In the absence of a theory of initial conditions, empiricists who use an incomplete dynamic model often take the first period of play as the default initial condition and restrict the dynamic model to subsequent periods. This is a bad approach for several reasons. First, the model remains incomplete, so empirical results cannot be used to predict complete dynamic paths including the first period. Second, in finite populations, the observed empirical frequency of play in the first period is the realization of a random variable, and hence does not reveal the relevant latent state variable. To assume that it does introduces possibly significant errors into the subsequent empirical analysis. Third, since the truth is at best only approximated by the dynamic model, increased accuracy of first period play will not necessarily yield the best T-period ahead predictions. Fourth, this approach is not applicable when an ex-ante prediction is needed for a new game.

2.2. Uniform Period 1 Initial Condition.

Common in reinforcement learning theories (Erev and Roth, 1998, Sarin and Vahid, 1999) is to assume that since players have no past experience and no deductive reasoning abilities, initial conditions can be ex ante predicted reasonably well by *the principle of insuffucent reason*: that is, all pure strategies are equally likely.

Erev and Roth (1998)[1] model initial propensities as a fixed point assigning equal propensity to each available action. Empirical evidence in their experiments appears favorable to that assumption. Equal propensities are similarly used by Capra et al (1999).

Whereas uniform initial propensities are not likely to be empirically accurate in many settings, this model has the advantage that it results in high initial variance, whereas more accurate models may result in insufficient diversity.

2.3. Mixtures of Heterogeneous Types.

A middle ground between the zero rationality of section 2.1 and the infinite rationality of deductive theories is a mixture model of heterogeneous types introduced by Stahl and Wilson (1994, 1995) and extended by HS99. This model includes random behavior, as well as boundedly rational based on conjectures by Stahl (1993) and Nagel (1995) that different behaviors were due to different levels of reasoning by a self-referential process starting with a uniform prior over other players' strategies. Hence a level-1 players would best-respond to a uniform prior, a level-2 player would best respond to a population of level-1 players, and so on.

Stahl and Wilson (1994, 1995) demonstrated robustness of level-n parameters over games, whereas Nagel et al (1999) show some fascinating regularities in level-n behavior over a diverse range of populations in diverse settings. The six types of behavior considered here for characterization of the initial distribution of choice are (1) random behavior, also known as level-0, (2) level-1 bounded rationality, (3) level-2 bounded rationality, (4) uniform Nash behavior[2], (5) maximax behavior[3], and (6) worldly behavior.

[1] Motivated by the same consideration, Roth and Erev (1995) present simulations with initial propensities chosen randomly from a uniform distribution over each player's pure strategies.

[2] By uniform Nash behavior we mean that each pure Nash equilibrium is equally likely. This is in accordance with HS99, which found no evidence for deductive selection in initial period play.

The last type of behavior mentioned, the worldly behavior, can be thought of as considering a convex combination of evidences, each of which underlies one of our previously defined behavioral types.

Each type of behavior was assigned a precision parameter, ν_t, which can be thought of as a parameter which determines how sensitive players of type t are to expected payoff differences between available actions. Through a likelihood function of a logit specification, both the precision parameters and the proportion of each type of behavior in the population, α_t, were determined. The parameters used for predictions of initial-period play in this paper (see Table 1) are based on out-of-sample estimation of 20 different one-shot games in HS99.

Given the 12 estimated parameters in Table 1, this model of initial play defines a probability distribution over the pure strategies, which is the derived estimate of the probability distribution from which the empirical sample was drawn. Hence, this estimated probability distribution can be used to forecast initial play out-of-sample.

2.4. Level-1 Initial Conditions.

Due to the large number of parameters (12), the model of HS99 may be too cumbersome to use. Alternatively, one could use the homogeneous level-1 model considered in HS99, since level-1 behavior was the predominant mode in the heterogeneous HS99 model. Recall from the previous section that the level-1 player noisily best responds to the prior that the population is composed of equal numbers of players choosing each available action. Stahl (1999) compared different population models of dynamics starting from level-1 initial conditions. Sonsino, Erev, and Gilat (1999) similarly applied level-1 initial conditions to a model of reinforcement learning that had player reinforcing rules.

2.5. Uniform Period 0.

---

[3] A maximax type is one who tends to choose the action that can potentially give him the highest payoff in the game (Haruvy, Stahl, and Wilson, 1999).

Stahl (1999) suggested imagining a "prior period 0" to which the principle of insufficient reason implies a uniform distribution, and then invoking the dynamic model for period 1 onward. This approach was demonstrably better than assuming a uniform distribution in period 1. The first period predictions are quite similar to those of the level-1 model, and have the advantage of requiring no additional parameters to complete the dynamic model. Therefore, we will examine this approach herein.

3. THE EXPERIMENT

3.1. The Game Environment

Consider a finite, symmetric, two-player game $G \equiv (N,A,U)$ in normal form, where $N \equiv \{1,2\}$ is the set of players, $A \equiv \{1, \ldots, J\}$ is the set of actions available to each player, and U is the $J\times J$ matrix of expected utility payoffs for the row player, and U', the transpose of U, is the payoff matrix for the column player. We focus on single population situations in which each player observes the frequency distribution of the past play of the other players in the population.

3.2. Experimental Design

An experiment session consisted of 25 players playing two runs of 12 periods each. In the first run , a single 3×3 symmetric game was played for 12 periods, and in the second run, a different 3×3 symmetric game was played for 12 periods. A "mean-matching" protocol was used. 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 \$2.00 per period of play. In other words, the token payoff for each period gave the percentage chance of winning \$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 ten-sided dice. A player won $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.

3.3. The experimental data

We select five games investigated in HS99, with properties that make them ideal for a thorough study of learning dynamics and equilibrium selection. These are games 1, 13, 14, 16, and 19 of HS99 (see figure 1). We further investigate two additional versions of game 13 (with 20 added to each cell, and with 20 subtracted from each cell) and an additional version of game 16 (with 20 added to each cell). Games 1, 14, and 19 begin with initial conditions far from uniform and very little movement is observed thereafter, with the exception of one run of game 14. The game 13 versions have long dynamic paths from initial conditions to the final outcome. The two versions of game 16 are characterized by an almost equal split in outcomes. Although the two equilibrium outcomes are equally observed, initial conditions fall in mainly one basin, resulting in a separatrix crossing. The games are shown in figure 1. Altogether we have five game 1 runs, seven game 13 runs (three versions), five game 14 runs, five of each game 16 version, and five game 19 runs, for a total of 32 runs.

4. Goodness of Fit of Initial Condition Predictions

In order to evaluate different initial conditions, we must decide on the appropriate measures of fit and apply those measure on the complete models. Unfortunately, there is no agreement in the literature on how to evaluate complete dynamic models. We group candidate measures into two categories: (1) one-period ahead, and (2) T-period ahead. The one-period-ahead category includes all the measures based on the one-period ahead predictions of the model - that is, the theoretical choice probabilities for period t conditional on the actual choices in period t-1. The standard log-likelihood measures belong to this category, and they are natural for parameter estimation because the dynamic models being fitted are Markovian (i.e. one-period ahead). The theoretical conditional choice probabilities can also be used to construct alternative goodness-of-fit measures such as the Root-Mean-Squared-Error (RMSE), and the Pearson Chi-square measure (PCS).

In contrast, T-period ahead measures are based on the unconditional choice probabilities for period T. Since closed-form solutions are rarely tractable, simulation techniques are used to generate a large pseudo-sample of dynamic paths from which various goodness-of-prediction measures can be computed. The stochastic nature of finite-population dynamics implies that T-period ahead predictions are probabilistic: yielding a non-degenerate probability distribution over the J-dimensional simplex (J being the number of actions in the game). Since experimental data sets rarely contain more than a dozen complete paths per game, small-sample problems arise. While one can construct an estimate of the likelihood density of the data paths, these measures computed for small sample sizes can be quite sensitive to the fine-structure of the T-period ahead probability distributions, and hence not reliable indicators of goodness-of-prediction. We suggest an alternative discrete-event measure that is reliable for small (but not too small) sample sizes. We divide the J-dimensional simplex into J+1 equal-sized regions, with each of J regions containing exactly one of the vertices of the simplex (see figure 2). The T-period-ahead probability of being in each region is computed from the pseudo-sample of dynamic paths. Thus, likelihood density on a continuum is converted into J+1 discrete-event probabilities; from these it is straightforward to compute likelihood measures.

5. Models of Adaptive Dynamics

We focus our study on four of the leading *population* learning models: (1) "Mental" Replicator (MR) dynamics, (2) Erev-Roth (1998) reinforcement learning, (3) Camerer-Ho experience weighted attraction (EWA) dynamics, and (4) logit best-reply with inertia and adaptive expectations (LBRIAE). These models where studied extensively in Stahl (1999), and the latter was the winner of a horse race using a large data set consisting of 5x5 and 3x3 symmetric normal-form games with population feedback. Since our focus is on the impact of initial conditions for a particular dynamic rather than comparisons across dynamics, we do not re-estimate the parameters of these dynamic models on the new data set, but instead use the estimated parameters from Stahl (1999).

The four dynamics investigated have different basins of attraction in the games investigated. Hence, the initial conditions are import. To illustrate this, the phase diagrams displaying MR and EWA dynamics and separatrix for games 16 and 19 are shown in Figure 3.

6. Results

We begin with the trivial observation that in terms of initial-period fit, the actual initials perform best; this establishes the benchmark for the best possible initial-period fit (Table 2). Far more interesting is the ranking of HS99 as the leading model in log-likelihood (LL), RMSE, and PCS for each game, with the uniform period 0 and level-1 approaches second and third, respectively[4] (). The uniform period 1 approach has worst initial-period fit. However, one should note that whereas HS99 entails 12 parameters, the level-1 model introduces one parameter only, whereas uniform period 0 and uniform period 1 require no additional parameters. Tradeoffs of parsimony and accuracy are often important considerations in the choice of a model.

[4] By RMSE, level-1 is slightly better than uniform period 0.

We proceed to examine one-period-ahead measures over periods 2-12. In Table 3, we present the log-likelihood for periods 2-12. By this measure, the actual initials perform dismally. The actual initials' standing does not improve looking at the PCS and RMSE measures.

Notwithstanding the dismal performance of actual initials, perhaps encouraging to studies of dynamics which apply one-period-ahead criteria is the fact that these criteria are fairly robust to the other alternative initial conditions considered here. With the exception of MR, differences in periods 2-12 likelihood within each dynamic over initial conditions remain below 2.05.[5] MR appears somewhat sensitive to initial conditions. Partly due to the robustness of one-period-ahead measures to *reasonable* initial conditions (i.e. in the neighborhood of level-1 behavior), horse-race rankings of the four models examined, by any of the one-period-ahead criteria, would also be robust to reasonable initial conditions. The winner of the horse race, LBRIAE, wins by a substantial margin taking each game separately as well as on overall performance.

Next, we turn our attention to T-period-ahead measures. Robustness over initial conditions is no longer the case. The model of HS99 is a clear winner in games 1 and 19 for all dynamics, but performs terribly in games 14 and 16 for all dynamics. It ends up losing in overall performance to the uniform-period-1 approach because of games 14 and 16.

To get to the bottom of this conundrum, we must compare the different games. We note that games 14 and 16 were characterized by separatrix crossings. Since the dynamic models we investigage are weak at capturing such crossings, disperse initials such as uniform period 1 initials prove helpful by generating some initials in the other basin. We conjecture that in games for which the separatrix is near the centroid or the noisy level-1 behavior, variability is essential for T-Period ahead goodness-of-prediction, and uniform period 1 initial conditions provide the greatest variability of all models of initial conditions under consideration.

[5] It is noteworthy that in terms of the period 2-12 log-likelihood, the best initial conditions for the Erev-Roth model are uniform period 1, as we should expect since the parameters were obtained by maximizing the period 1-12 likelihood under this initial condition assumption.

On the other hand, uniform period 0 initial conditions are quite reliable at predicting initial play. Can these apparently contradictory findings be reconciled? When uniform period 1 initial conditions are used for a T-period ahead measure, the first period of play is simulated, so although the a priori probability of each choice is 1/J, the realized first period choices have a multinomial distribution, and this diversity is crucial for T-period ahead prediction in games like 16. In contrast, when uniform period 0 initial conditions are used, the fictitious period 0 is not simulated, and so lacks diversity. If it were simulated, the T-period simulated path would be virtually the same as a T+1 period simulated path from uniform period 1 initial conditions. This interpretation of our findings, suggest that the overall best model of initial conditions would be uniform period 0 *with diversity*.[6]

7. Conclusions

For one-period-ahead measures of goodness-of-fit the winners are, HS99, the uniform-period-0 approach, and level-1 in that order. For the T-period-ahead discrete-event measure, these same initial condition models perform well except for games 14 and 16, where uniform-period-1 wins, due to the added diversity of simulating a T-period path from uniform initial conditions.

In contrast to the one-period-ahead measures, ranking of dynamics using T-period ahead measures critically depends on the choice of initial conditions, with EWA winning for all but HS99 initials, and LBRIAE a close second.

[6] To provide a rigorous test of this conjecture, we would need to estimate the parameters of a model that explicitly incorporates period 0 diversity. Since we want to focus this paper on initial conditions and not parameter estimation, we leave this task for the future.

References

Camerer, C. F., and T-H Ho (1999), 'Experience-Weighted Attraction Learning in Normal-form Games,' *Econometrica* 67, 827-874.

Capra, M., J. Goeree, R. Gomez, and C. Holt (1999), "Anomalous Behavior in a Traveler's Dilemma?" *American Economic Review*, 89(3), 678-690.

Costa-Gomez, M., Crawford, V., and B. Broseta (1998) "Cognition and behavior in normal-form games: An experimental study," working paper, University of California, San Diego.

Erev, I. And A. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed Strategy Equilibria," *American Economic Review* 88, 848-881.

Harsany, J. (1995), "A New Theory of Equilibrium Selection for Games with Incomplete Information," *Games and Economic Behavior* 10, 318-332

Harsanyi, J., and R. Selten (1988*), A General Theory of Equilibrium Selection in Games*, M.I.T. Press.

Haruvy, E., and D. O. Stahl (1999) "Empirical Tests of Equilibrium Selection based on Player Heterogeneity," Revised and resubmitted in *Econometrica*.

Hofbauer, J., and K. Sigmund (1988), "The Theory of Evolution and Dynamical Systems," Cambridge University Press.

Nagel, R. (1995). "Unraveling in Guessing Games: An experimental study," *American Economic Review* 85, 1313-1326.

Nagel, R., A. Bosch, A. Satorra, and J. García (1999), "One, two, (three), infinity: Newspaper and Lab Beauty-Contest Experiments" Universitat Pompeu Fabra working paper.

Roth, A. and I. Erev (1995), "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior* **8**, 164-212.

Sonsino, D., I. Erev, and S. Gilat (1999), "On Rationality, Learning, and Zero-Sum Betting – An Experimental Study of the No-Betting Conjecture." Working paper, Technion, Israel.

Stahl, D. (1993). "Evolution of Smart$_n$ Players," *Games and Economic Behavior* 5, 604-617.

Stahl, D. O. (1996) "Boundedly Rational Rule Learning in a Guessing Game," *Games and Economic Behavior* **16**, 303-330.

Stahl, D. O. (1999) "A Horse Race Among Action Reinforcement Learning Models," mimeo.

Sarin, R. and F. Vahid (1999) “Payoff Assessments without Probabilities: A Simple Dynamic Model of Choice,” *Games and Economic Behavior*, 28, 294-309.

Stahl, D. and P. Wilson (1994), "Experimental Evidence of Players' Models of Other Players," *Journal of Economic Behavior and Organizations*, **25**, 309-327.

Stahl, D. and P. Wilson (1995), "On Players Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, **10**, 218-254.

Table 1

Parameter	Estimate
ν_1	0.560
ν_2	0.480
ν_{max}	0.331
ν_{NE}	0.902
ν_W	0.248
ε_w	0.845
μ_w	0.045

Parameter	Estimate
α_0	0.079
α_1	0.310
α_2	0.039
α_{max}	0.127
α_{NE}	0.023
α_W	0.422

Key: . The parameters ν_1 and ν_2 are the precision parameters for level-1 and level-2, respectively. The precision parameter ν_{NE} corresponds to the Nash type and ν_w to the worldly type. The parameter μ_w is the precision parameter for the boundedly rational types in the worldly evidence; ε_w is the mixture parameter in the worldly type's prior; and α_t is the proportion of population using rule t, where $t \in \{0,1,2\}$ denotes the level-t rule, t = max denotes the maximax rule, t = NE denotes the Nash evidence, and t = W denotes the worldly rule.

Table 2. Goodness of Fit -- Initial Period

Log Likelihood

Game	Uniform P(1)[7]	Level-1	Uniform P(0): M. Rep.	Uniform P(0): RE	Uniform P(0): EWA	Uniform P(0): LBRIAE	HS99	Actual
1	-27.47	-24.04	-23.54	-24.56	-23.62	-22.62	-23.34	-20.55
13	-27.31	-23.46	-23.90	-24.36	-24.06	-23.44	-23.42	-20.91
14	-27.25	-7.51	-9.47	-9.44	-9.88	-14.64	-6.23	-5.02
16	-27.14	-23.57	-24.06	-23.03	-24.20	-22.94	-23.72	-21.68
19	-27.25	-32.86	-27.24	-29.29	-27.74	-24.27	-25.09	-22.68
Total	-136.40	-111.44	-108.21	-110.69	-109.51	-107.91	-101.80	-90.84

PCS

Game	Uniform P(1)	Level-1	Uniform P(0): M. Rep.	Uniform P(0): RE	Uniform P(0): EWA	Uniform P(0): LBRIAE	HS99	Actual
1	13.52	7.53	5.24	6.517	5.270	3.347	7.64	0
13	11.99	4.25	4.95	5.846	5.240	4.308	5.75	0
14	42.79	3.03	5.79	5.711	6.382	14.032	1.43	0
16	11.07	3.32	4.28	2.308	4.559	2.168	5.81	0
19	8.66	70.58	14.45	23.550	17.060	2.737	5.30	0
Total	17.61	17.74	6.94	8.787	7.702	5.318	5.19	0

RMSE

Game	Uniform P(1)	Level-1	Uniform P(0): M. Rep.	Uniform P(0): RE	Uniform P(0): EWA	Uniform P(0): LBRIAE	HS99	Actual
1	0.245	0.133	0.134	0.175	0.139	0.116	0.134	0
13	0.231	0.133	0.148	0.164	0.153	0.129	0.131	0
14	0.438	9.5E-02	0.151	0.152	0.161	0.262	0.048	0
16	0.223	0.124	0.142	0.094	0.147	0.095	0.112	0
19	0.196	0.170	0.144	0.170	0.147	0.101	0.134	0
Total	0.281	0.133	0.144	0.154	0.150	0.154	0.116	0

[7] *Note*: Log-likelihood for the initial period under the uniform period 1 approach differs over games due to the different number of players (24-25) in different runs.

Table 3 Goodness of Fit -- One period-ahead measures -- Periods 2-12

Mental Replicator:

Init. Cond's:	LL(2-12):	RMSE(2-12):	PCS(2-12):
Uniform P1			
1	-110.328	0.096	39.203
13	-203.938	0.134	49.053
14	-136.119	0.087	21.018
16	-158.228	0.105	42.910
19	-126.308	0.099	34.228
Total	-734.92	0.105	37.282
Level-1			
1	-107.045	0.080	36.387
13	-205.670	0.140	52.625
14	-136.724	0.092	26.084
16	-155.742	0.094	39.256
19	-122.402	0.079	29.031
Total	-727.583	0.099	36.677
Uniform P0			
1	-107.239	0.081	36.420
13	-205.449	0.139	52.165
14	-136.401	0.090	24.970
16	-155.906	0.095	39.453
19	-122.667	0.081	29.389
Total	-727.662	0.099	36.480
HS99			
1	-103.592	0.070	31.796
13	-208.691	0.150	59.049
14	-138.118	0.097	30.922
16	-152.997	0.086	36.741
19	-120.866	0.073	26.506
Total	-724.264	0.010	37.003
Actual			
1	-105.41	0.077	33.888

13	-206.023	0.141	53.721
14	-141.654	0.106	44.302
16	-153.686	0.086	36.800
19	-122.301	0.081	28.307
Total	-729.075	0.101	39.404

Erev-Roth:

Init. Cond's:	**LL(2-12):**	**RMSE(2-12):**	**PCS(2-12):**
Uniform P1			
1	-101.678	0.068	30.778
13	-199.828	0.120	45.835
14	-134.925	0.077	23.682
16	-151.866	0.078	38.234
19	-120.683	0.068	36.753
Total	-708.981	0.084	35.056
Level-1			
1	-101.77	0.065	32.731
13	-200.32	0.122	47.265
14	-135.92	0.080	27.632
16	-151.72	0.077	38.095
19	-120.37	0.065	36.540
Total	-710.10	0.084	36.453
Uniform P0			
1	-102.429	0.070	31.962
13	-200.069	0.121	46.161
14	-135.062	0.079	24.276
16	-151.564	0.076	38.225
19	-121.788	0.070	40.442
Total	-710.913	0.085	36.213
HS99			
1	-100.832	0.062	31.360
13	-200.919	0.123	48.972
14	-136.162	0.080	28.707
16	-151.561	0.076	38.036

19	-120.131	0.064	36.093
Total	-709.606	0.084	36.634
Actual			
1	-101.138	0.065	31.535
13	-200.811	0.123	48.694
14	-136.188	0.080	28.792
16	-151.604	0.076	38.034
19	-120.357	0.066	36.324
Total	-710.097	0.085	36.676

EWA Dynamics:

Init. Cond's:	LL(2-12):	RMSE(2-12):	PCS(2-12):
Uniform P1			
1	-98.401	0.056	37.858
13	-194.412	0.105	34.224
14	-132.743	0.073	20.320
16	-157.465	0.098	40.941
19	-117.754	0.059	37.143
Total	-700.774	0.081	34.097
Level-1			
1	-98.63	0.056	38.366
13	-195.09	0.106	34.641
14	-133.47	0.074	21.053
16	-156.69	0.098	40.758
19	-117.92	0.059	37.520
Total	-701.81	0.081	34.468
Uniform P0			
1	-98.568	0.056	37.858
13	-194.924	0.105	34.224
14	-133.26	0.073	20.320
16	-156.829	0.098	40.941
19	-117.885	0.059	37.143
Total	-701.466	0.081	34.097
HS99			

1	-97.9523	0.055	37.320
13	-195.943	0.110	36.858
14	-133.848	0.075	22.361
16	-155.273	0.093	38.900
19	-117.616	0.059	36.207
Total	-700.633	0.081	34.329
Actual			
1	-98.135	0.055	37.225
13	-195.820	0.108	36.953
14	-135.218	0.079	26.979
16	-155.872	0.095	39.683
19	-117.630	0.059	35.671
Total	-702.674	0.082	35.302

LBRIAE

Init. Cond's:	LL(2-12):	RMSE(2-12):	PCS(2-12):
1	-95.520	0.058	22.508
13	-190.892	0.088	27.151
14	-134.287	0.072	23.225
16	-149.752	0.075	29.497
19	-115.203	0.057	20.905
Total	-685.664	0.071	24.657

Table 4: Goodness of Fit – T-period-ahead Discrete Event Measure

Period 12 Hit Rates and LL Measure by Model and Game

	A	B	C	Other	
1	5	0	0	0	5
13	0	0	6	1	7
14	0	4	1	0	5
16	5	5	0	0	10
19	0	5	0	0	5

Mental Replicator

Uniform Period 1:

Game	A region	B region	C region	Central region	LL
1	0.515	0.458	0.000	0.027	-0.288
13	0.009	0.000	0.978	0.013	-0.008
14	0.000	0.920	0.074	0.006	-0.255
16	0.725	0.274	0.000	0.001	-0.351
19	0.630	0.347	0.000	0.023	-0.460
					-1.363

Level-1:

Game	A region	B region	C region	Central region	LL
1	0.518	0.451	0.000	0.031	-0.286
13	0.008	0.000	0.979	0.013	-0.008
14	0.000	0.982	0.014	0.005	-0.380
16	0.881	0.118	0.000	0.001	-0.491
19	0.657	0.316	0.000	0.027	-0.501
					-1.665

Uniform Period 0:

Game	A region	B region	C region	Central region	LL
1	0.489	0.480	0.000	0.031	-0.311
13	0.008	0.000	0.980	0.013	-0.008
14	0.000	0.977	0.018	0.005	-0.358
16	0.868	0.131	0.000	0.001	-0.472
19	0.705	0.271	0.000	0.025	-0.567
					-1.716

HS99:

Game	A region	B region	C region	Central region	LL
1	0.980	0.014	0.000	0.006	-0.009
13	0.007	0.000	0.981	0.012	-0.007
14	0.000	0.994	0.004	0.003	-0.489

16	0.947	0.052	0.000	0.001	-0.653
19	0.128	0.846	0.000	0.026	-0.072
					-1.230

EWA Dynamics

Uniform Period 1:

Game	**A region**	**B region**	**C region**	**Central region**	**LL**
1	0.796	0.203	0.001	0.000	-0.099
13	0.010	0.002	0.975	0.013	-0.009
14	0.000	0.889	0.107	0.005	-0.235
16	0.706	0.283	0.000	0.010	-0.350
19	0.277	0.723	0.000	0.000	-0.141
					-0.834

Level-1:

Game	**A region**	**B region**	**C region**	**Central region**	**LL**
1	0.959	0.041	0.000	0.000	-0.018
13	0.004	0.000	0.985	0.012	-0.006
14	0.000	0.963	0.033	0.004	-0.310
16	0.850	0.139	0.000	0.011	-0.463
19	0.075	0.925	0.000	0.000	-0.034
					-0.831

Uniform Period 0:

Game	**A region**	**B region**	**C region**	**Central region**	**LL**
1	0.923	0.077	0.000	0.000	-0.035
13	0.004	0.000	0.985	0.012	-0.006
14	0.000	0.947	0.049	0.005	-0.282
16	0.825	0.165	0.000	0.011	-0.433
19	0.147	0.853	0.000	0.000	-0.069
					-0.825

HS99:

Game	**A region**	**B region**	**C region**	**Central region**	**LL**
1	1.000	0.000	0.000	0.000	0.000
13	0.004	0.000	0.984	0.012	-0.006
14	0.000	0.975	0.021	0.004	-0.344
16	0.910	0.081	0.000	0.009	-0.567
19	0.001	0.999	0.000	0.000	0.000
					-0.917

Erev-Roth:

Uniform Period 1:

Game	**A region**	**B region**	**C region**	**Central region**	**LL**
1	0.555	0.418	0.002	0.025	-0.256
13	0.105	0.008	0.862	0.026	-0.055

14	0.000	0.884	0.110	0.007	-0.235
16	0.729	0.271	0.000	0.000	-0.352
19	0.520	0.449	0.000	0.031	-0.348
					-1.247

Level-1:

Game	A region	B region	C region	Central region	LL
1	0.772	0.206	0.000	0.022	-0.113
13	0.039	0.000	0.941	0.021	-0.023
14	0.000	0.994	0.005	0.001	-0.466
16	0.885	0.115	0.000	0.000	-0.497
19	0.259	0.712	0.000	0.029	-0.147
					-1.245

Uniform Period 0:

Game	A region	B region	C region	Central region	LL
1	0.539	0.430	0.000	0.031	-0.268
13	0.029	0.000	0.957	0.014	-0.016
14	0.000	0.976	0.020	0.004	-0.348
16	0.976	0.024	0.000	0.000	-0.813
19	0.539	0.427	0.000	0.034	-0.370
					-1.816

HS99:

Game	A region	B region	C region	Central region	LL
1	0.974	0.020	0.000	0.007	-0.012
13	0.061	0.000	0.913	0.026	-0.034
14	0.000	0.998	0.001	0.001	-0.571
16	0.922	0.078	0.000	0.000	-0.572
19	0.095	0.888	0.000	0.017	-0.051
					-1.241

LBRIAE

Game	A region	B region	C region	Central region	LL
1	0.925	0.073	0.000	0.002	-0.034
13	0.013	0.000	0.958	0.029	-0.016
14	0.000	0.977	0.022	0.002	-0.340
16	0.826	0.172	0.000	0.002	-0.424
19	0.072	0.925	0.000	0.003	-0.034
					-0.848

Figure 1. The Games

Game 1

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

Game 19

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

Game 14

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

Game 13

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

Game 13 (-20)

	A	B	C
A	40	40	10
B	10	50	0
C	50	5	15

Game 13 (+20)

	A	B	C
A	80	80	50
B	50	90	40
C	90	45	55

Game 16

	A	B	C
A	20	0	60
B	0	60	0
C	10	25	25

Game 16 (+20)

	A	B	C
A	40	20	80
B	20	80	20
C	30	45	45

Figure 2 The discrete-event measure

Dividing the J-dimensional simplex into J+1 equal-sized regions, with each of the J regions containing exactly one of the vertices of the simplex

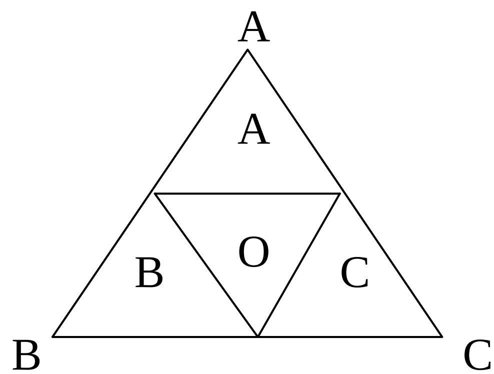

Figure 3

Game 16 Mental Replicator

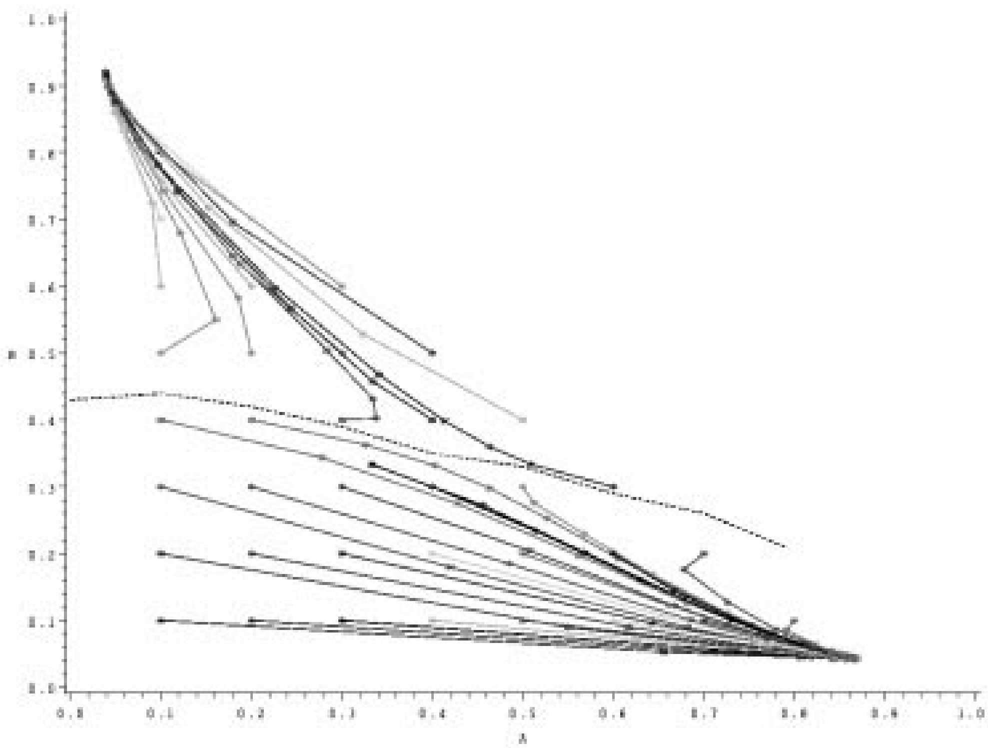

Game 16 EWA

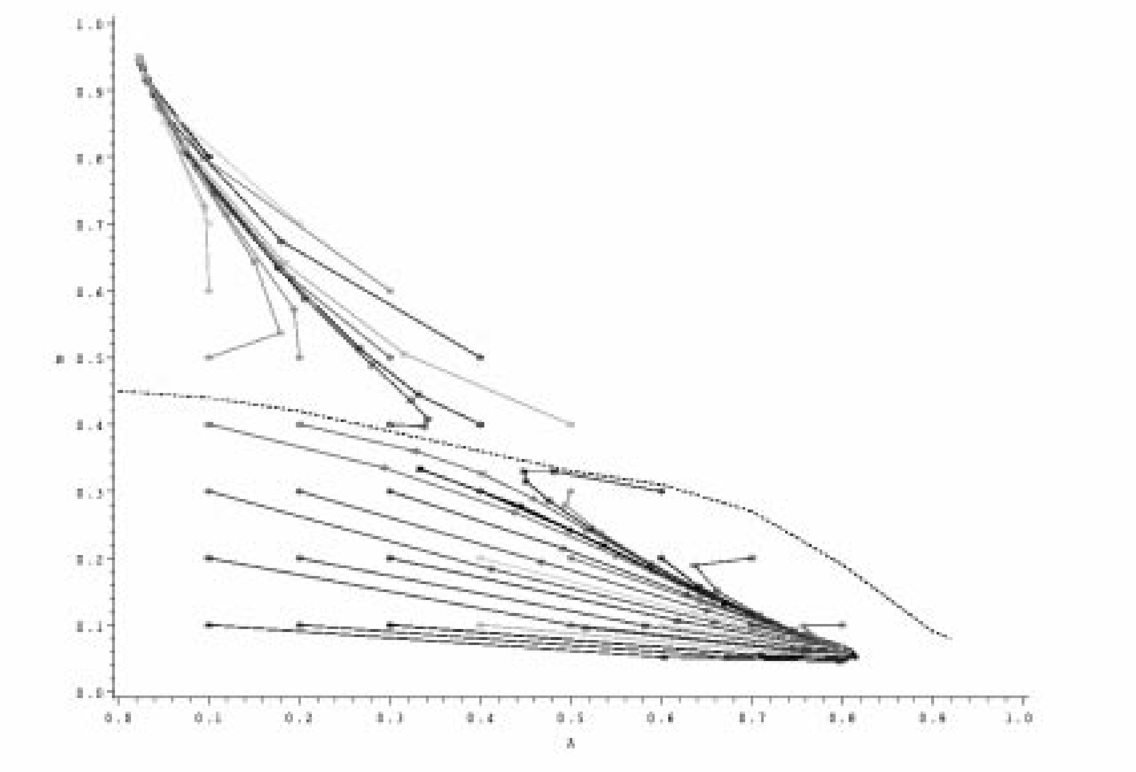

Game 19 Mental Replicaotr

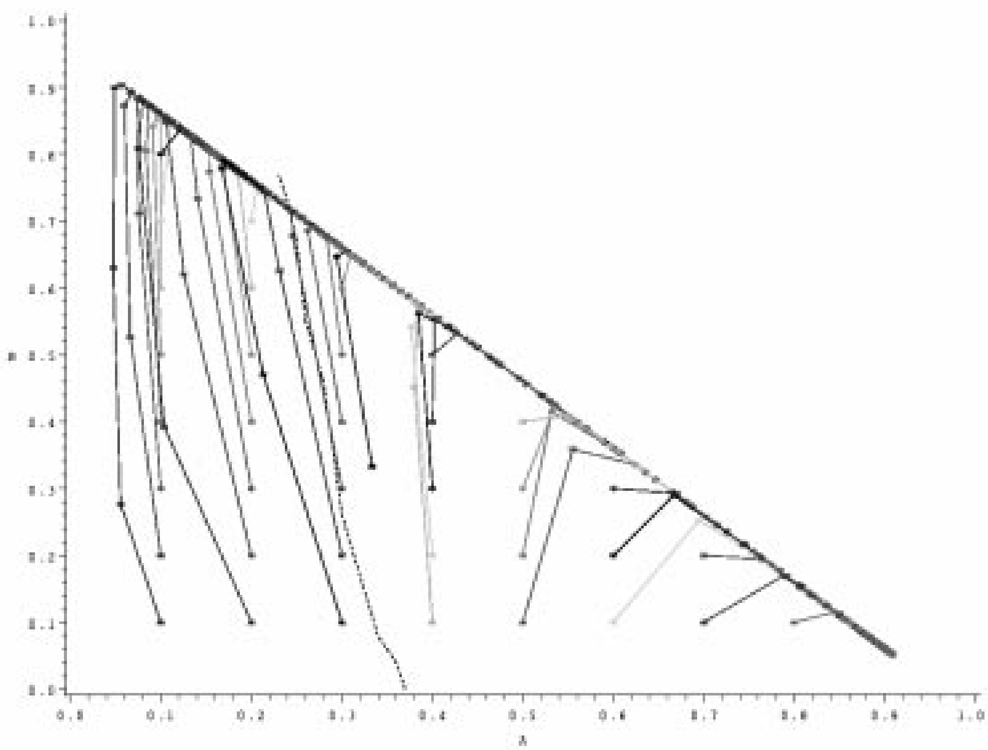

Game 19 EWA