

Equilibrium Convergence in Normal Form Games

Nicole Bouchez and Daniel Friedman
Economics Department University of California, Santa Cruz

November 2000

In this chapter, we examine convergence behavior in simple bimatrix games. We classify the possible types of simple games, pick interesting examples of each type, and summarize convergence behavior under various information and player matching protocols. See Friedman (1996), Cheung and Friedman (1997) and Bouchez (1998) for more complete descriptions of the experiments.

We begin with normal form games that have only two alternative strategies and a single symmetric population of players. These games are defined by a 2×2 matrix $A = ((a_{ij}))$ specifying the payoff to any player choosing strategy i when the opponent chooses j .

Evolutionary game theory predicts that the direction of change in the fraction $p \in (0,1)$ of players choosing the first strategy is given by the sign of the payoff differential

$D(p) = (1, -1)A(p, 1-p)' = (1-p)a - pb$ where $a = a_{12} - a_{22}$ and $b = a_{21} - a_{11}$. When $D(p)$ is positive (i.e., the first strategy has the higher payoff) then $p < 1$ increases and the fraction $1-p$ of players choosing the alternative strategy $i=2$ decreases; the opposite is true when $D(p)$ is negative. The graph of $D(p)$ is a straight line with intercept a at $p = 0$ and value $-b$ at $p = 1$.

Thus (apart from the degenerate case $a=b=0$ in which a player is always indifferent between her two actions) each payoff matrix falls into one of three qualitatively different types as

shown in Figure 1. In accordance with this classification scheme, we used single and two population games of all three types.

The next most complicated case is a single population of strategically identical players with three alternative actions. Here the payoff matrix \mathbf{A} is 3×3 and the current state is a point in two-dimensional simplex $S = \{(p, q, 1-p-q) \in \mathbf{R}^3: p, q \geq 0, p+q \leq 1\}$. The classification of matrices becomes more complex as the edges of the simplex retain all three possibilities and the interior can be a sink, source, saddle or center. We consider here only a version of the "Hawk-Dove-Bourgeois" (HDB) game, as illustrated in Figure 2.

I. LABORATORY PROCEDURES AND TREATMENTS

The experiments consist of 60-120 minute laboratory sessions with 6 to 24 undergraduate subjects. Population size varies from 8 to 16 in the results presented here; Friedman (1996) finds strategic behavior contrary to the evolutionary assumption with population sizes of 6 and smaller. After instruction and a few practice periods, each session consists of 60-200 periods broken up into runs of 10 to 16 periods. Over 90% of the subjects earned between \$8 and \$32 per session.

The treatments used were *random pairwise* (RP) and *mean matching* (MM) matching protocols, and the amount of historical information that appears on each player's screen (Hist/No Hist).

Under *random pairwise* (RP) the player has a single opponent each period chosen randomly and independently from the set of possible opponents. Players view their own (but not the opponent's) payoff matrix, and type "a" or "b" at the keyboard to indicate the choice of the

first or second strategy. Under RP matching for payoff matrix A , a player's choice of strategy $i = 1$ or 2 , gives *expected* payoff $(2-i, i-1)A(p, 1-p)$ when the fraction of potential opponents choosing strategy 1 is p . However, his *actual* payoff depends on the action taken by his actual opponent, and so has some variance around its expectation. The variance is eliminated in the alternative matching procedure, called *mean matching* (MM). Here each player is matched once against each possible opponent in each round and receives the average (mean) payoff over all his matches.

The other major treatment in our experiments is the amount of historical information that appears in the upper left box on each player's screen. In the minimum level, *No Hist*, the player receives no historical information other than what she could tabulate herself: her own action and actual payoff in previous periods. In the other level, *Hist*, the box also displays the full distribution of choices in the opponent population in previous periods.

All treatments are held constant within a run to test for convergence. Runs are separated by obvious changes in the player population and/or the payoff matrix, and the history box is erased at the beginning of a new run.

II. RESULTS

We have collected more than 300 such runs and used various statistical tests as well as summary graphs to study convergence properties. Figures 3-5 illustrate behavior observed in different sorts of environments. The main findings, presented more fully in Friedman (1996), can be summarized as follows.

1) Some behavioral equilibrium (BE) is typically achieved by the second half of a 10 to 16 period run. The operational definition of BE is that strategy selection is almost constant in

each population in a given run (or half-run). Almost constant means that the mean absolute deviation from the median number of players choosing a given strategy is less than 1 player ("tight") or less than 2 players ("loose"). Overall, we observe tight BE in over 70% of second half-runs and loose BE in over 98% of second half-runs. Tight BE was achieved most reliably in type 3 games (over 95% of all half-runs). By contrast, type 1 games achieved tight BE in only 55% of half runs, but achieved loose BE in 96%.

2) BE typically coincides with NE, especially with those (called EE) that evolutionary game theory identifies as stable. We operationalize NE by replacing the median number of players by the NE number. In second half-runs, for example, about 79% of the loose BE are loose NE, and 84% of those are loose EE. The only notable exception to this conclusion is that in type 2 runs the BE sometimes coincided with the non-EE mixed NE. A closer look at the graphs suggests that many of these cases actually represent slow divergence from the mixed NE, and many of the half-runs deemed BE but not NE seem to represent slow or incomplete convergence to an EE (a pure NE.)

3) Convergence to BE is faster in the mean-matching (MM) than in the random-pairwise (RP) treatment, and faster in the *Hist* treatment than in the *No Hist* treatment. In particular, the slow and incomplete convergence observed in type 1 games arises mainly in RP matching protocol and No Hist runs. The results from type 2 single population games and all two population games support the same conclusion. There is, however, an interesting exception. The few instances of non-convergence in type 3 games arise more often under MM than under RP.

4) Individual behavior at a mixed strategy BE is better explained by idiosyncratic "purification" strategies than by identical individual mixed strategies. In particular, in the

simplest type 1 game, Hawk-Dove, we see persistent heterogeneity in which some players consistently pick the first ("Hawk") strategy and others consistently pick the other ("Dove") strategy.

5) "Hawk-Dove-Bourgeois" is a 1-population 3-action game with a triangular state space and with one corner NE (an EE) with target area b^2 and one edge NE (not an EE) with target area $2b^2$. Only one session was explored in Friedman 1996. Additional data has been collected (Bouchez 1998) and the results summarized here are for the combined data. Loose (tight) convergence was found to some BE in 41 (7) of 46 half-runs, loose (tight) convergence to the EE in 8 (3) half-runs, and no loose or tight convergence was found to the edge NE despite its larger area. The data are sparse but consistent with evolutionary game theory.

III. DISCUSSION

For all three types of one dimensional games and their two dimensional analogues, the states reliably achieve a loose behavioral equilibrium (BE) even within the first half-run of 5 periods. Most of the loose BE are also tight BE, the main exceptions occurring in two dimensional games with unique Nash equilibria (NE). Most BE coincide with NE, and most of the observed NE are indeed evolutionary equilibria (EE). In general, the "evolutionary" treatments of mean-matching (MM) and feedback (Hist) appear to improve convergence to EE. Thus, the main tendencies of the data are consistent with evolutionary game theory.

The exceptions or boundaries to these main tendencies may be of special interest. Friedman (1996) shows that when group size is smaller than 6, players much more often appear willing to sacrifice current personal payoffs to increase group payoffs (and perhaps own later payoffs). Cooperative behavior (foregoing the dominant strategy) is sometimes observed in

type 3 prisoner's dilemma sessions that have runs splitting the players into groups of size 2 or 4, and it is especially prevalent in the runs with the smaller groups. Such behavior is notably less frequent in sessions where the minimum group size remains above 6.

Perhaps the most surprising finding concerns another boundary for evolutionary game theory. An influential branch of the theory (Kandori, Mailath and Rob, 1993, and Young, 1993) argues that in simple coordination (type 2) games with two pure strategy (corner) NE = EE and one interior NE, the "risk-dominant" corner EE is most likely to be observed because it has the larger basin of attraction, and indeed that *only* the risk-dominant EE will be observed in the relevant limiting case. Friedman (1996) shows that the data reviewed in this chapter support the contrary theoretical view of Bergin and Lipman (1995) that one can bias convergence towards the other ("payoff-dominant") EE by increasing the potential gains to cooperation. In some applications, evolutionary game theory may have to be supplemented by a theory of trembles (or "mutations") that allows for forward-looking attempts to influence others' behavior.

References

- Bergin, J. and Lipman, B. L. (1995). 'Evolution with state-dependent mutations.' Draft manuscript. Queen's University Department of Economics.
- Bouchez, Nicole (1998) 'A Study and Application of Learning in Population Games.' Draft manuscript. University of California, Santa Cruz.
- Cheung, Yin-Wong and Friedman, Daniel (1997). 'Individual learning in normal form games: some laboratory results.' *Games and Economic Behavior*, 19, 46-76.

Friedman, Daniel (1996). 'Equilibrium in Evolutionary Games : Some Experimental Results.'

The Economic Journal, vol. 106 (434), pp. 1-25, January.

Kandori, Michihiro, Mailath, George and Rob, Rafael (1993). 'Learning, mutations, and long

run equilibria in games.' *Econometrica* vol. 61 (1), pp. 29-56, January.

Young, H. Peyton (1993). 'The evolution of conventions.' *Econometrica* vol. 61 (1), pp. 57-


84, January.

Weibull, Jorgen W. (1995). *Evolutionary Game Theory*. Cambridge, Massachusetts: MIT


Press.

Figure 1: The three generic types of 2x2 symmetric games. The 2x2 matrix $A=(a_{ij})$ specifies the payoff to any player choosing strategy i when the opponent chooses j . The direction of change in the fraction p of players choosing the first strategy is the sign of the payoff differential $D(p) = (1-p)A(p,1-p) - pA(1-p,p) = (1-p)a - pb$ where $a=a_{12}-a_{22}$ and $b=a_{21}-a_{11}$. The slope of $D(p)$ and location of the root $p^*=a/(a+b)$ of $D(p)=0$ determine the type of game A .


Type 1: $a, b > 0$
 We have global convergence to $p^*=a/(a+b)$ over $[0,1]$ because D is downward sloping so p increases (decreases) whenever it is below (above) p^* . Hence p^* is the unique evolutionary equilibrium (EE). It is also the unique Nash equilibrium (NE).




Type 2: $a, b < 0$
 $p^* = a/(a+b)$ separates the basins of convergence of the two evolutionary equilibria $p = 0$ and $p = 1$. The NE include p^* as well.



Type 3: If a and b have opposite signs, then $D(p)$ lies everywhere above (or everywhere below) the p -axis, i.e. the first pure strategy $p = 1$ (or the second pure strategy $p = 0$) is dominant. The dominant strategy is the unique EE (and the unique NE).



or






Figure 2: A 3x3 symmetric game. This version of the Hawk-Dove-Bourgeois game has a corner Nash equilibrium (NE) at $(p, q) = (0, 0)$ and an edge NE at $(2/3, 1/3)$. Under standard evolutionary dynamics, the corner NE is an evolutionary equilibrium (EE) and the edge NE is a saddle point. The equations $d\mathbf{p}/dt = (1, 0, 0)' \mathbf{A} \mathbf{s} - (1/3, 1/3, 1/3)' \mathbf{A} \mathbf{s}$ and $d\mathbf{q}/dt = (0, 1, 0)' \mathbf{A} \mathbf{s} - (1/3, 1/3, 1/3)' \mathbf{A} \mathbf{s}$ define one simple version of the dynamics in the interior of the simplex; see Weibull (1995) for an introduction to replicator and other standard dynamics.

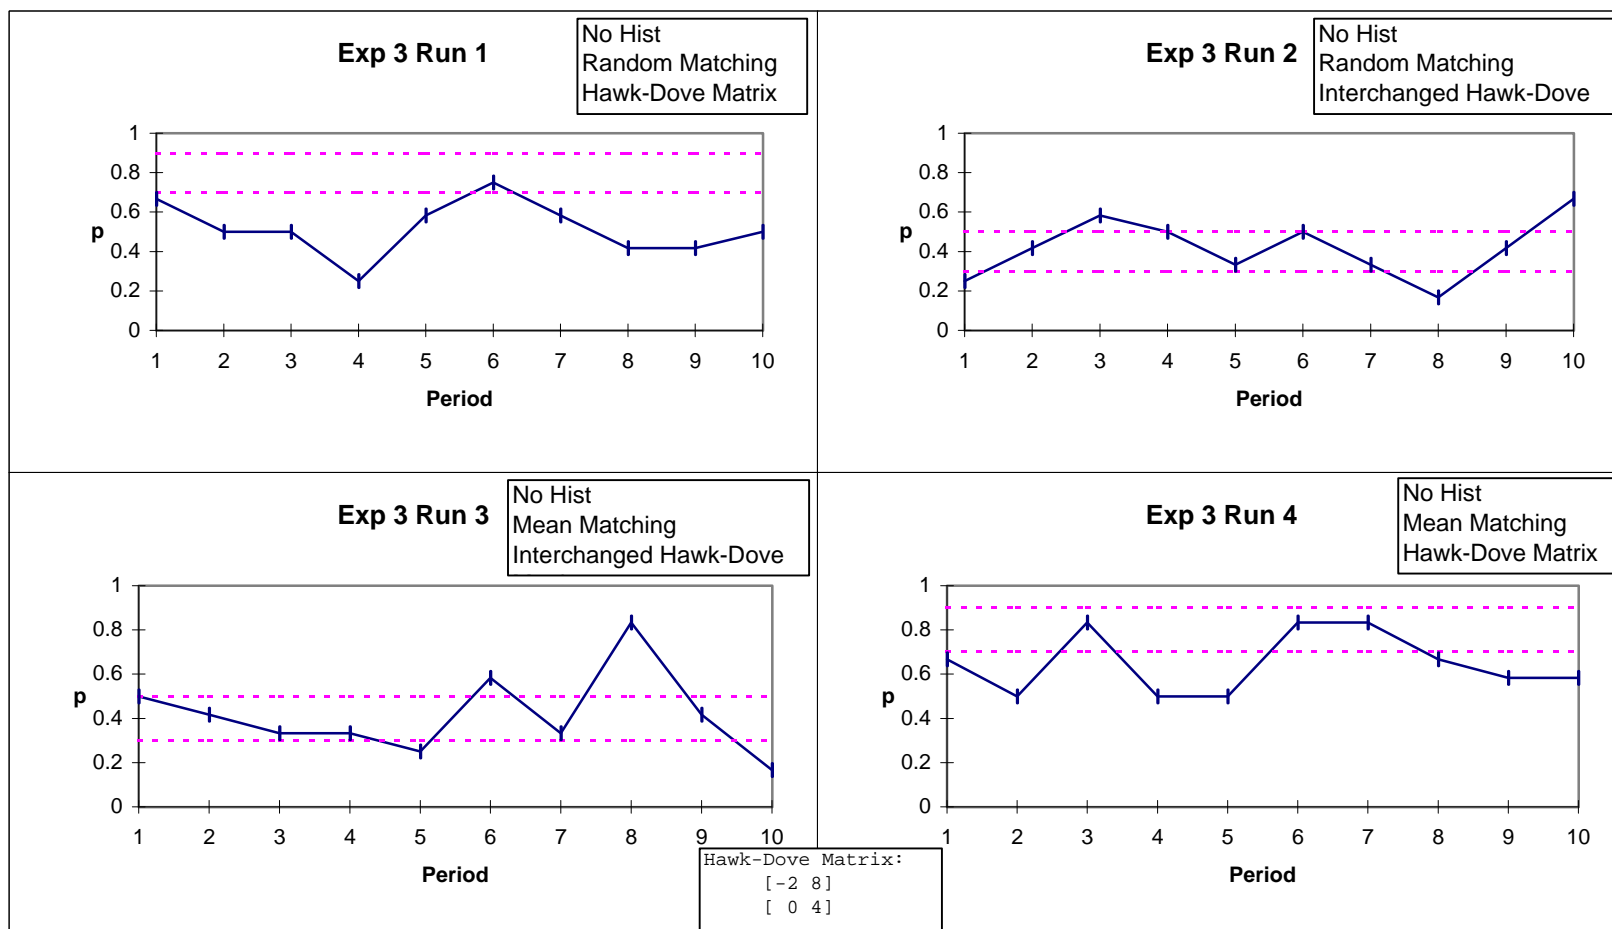


Figure 3: The time path of p_t in the first four runs of Exp3, the first usable session. The type 1 payoff matrix here has unique mixed NE $p^* = 2/3 = 8/12$. That is, in NE 8 of 12 players choose the first strategy (or 4 of 12 when the matrix rows and columns are interchanged as in runs 2 and 3). The graphs show a tolerance of 1 player in the band around NE. The time paths in the first four runs suggest that the NE attracts states p_t outside the tolerance band $p^* \pm 1/12$, but there is considerable behavioral noise, the path is inside the band in only about 50% of the periods.

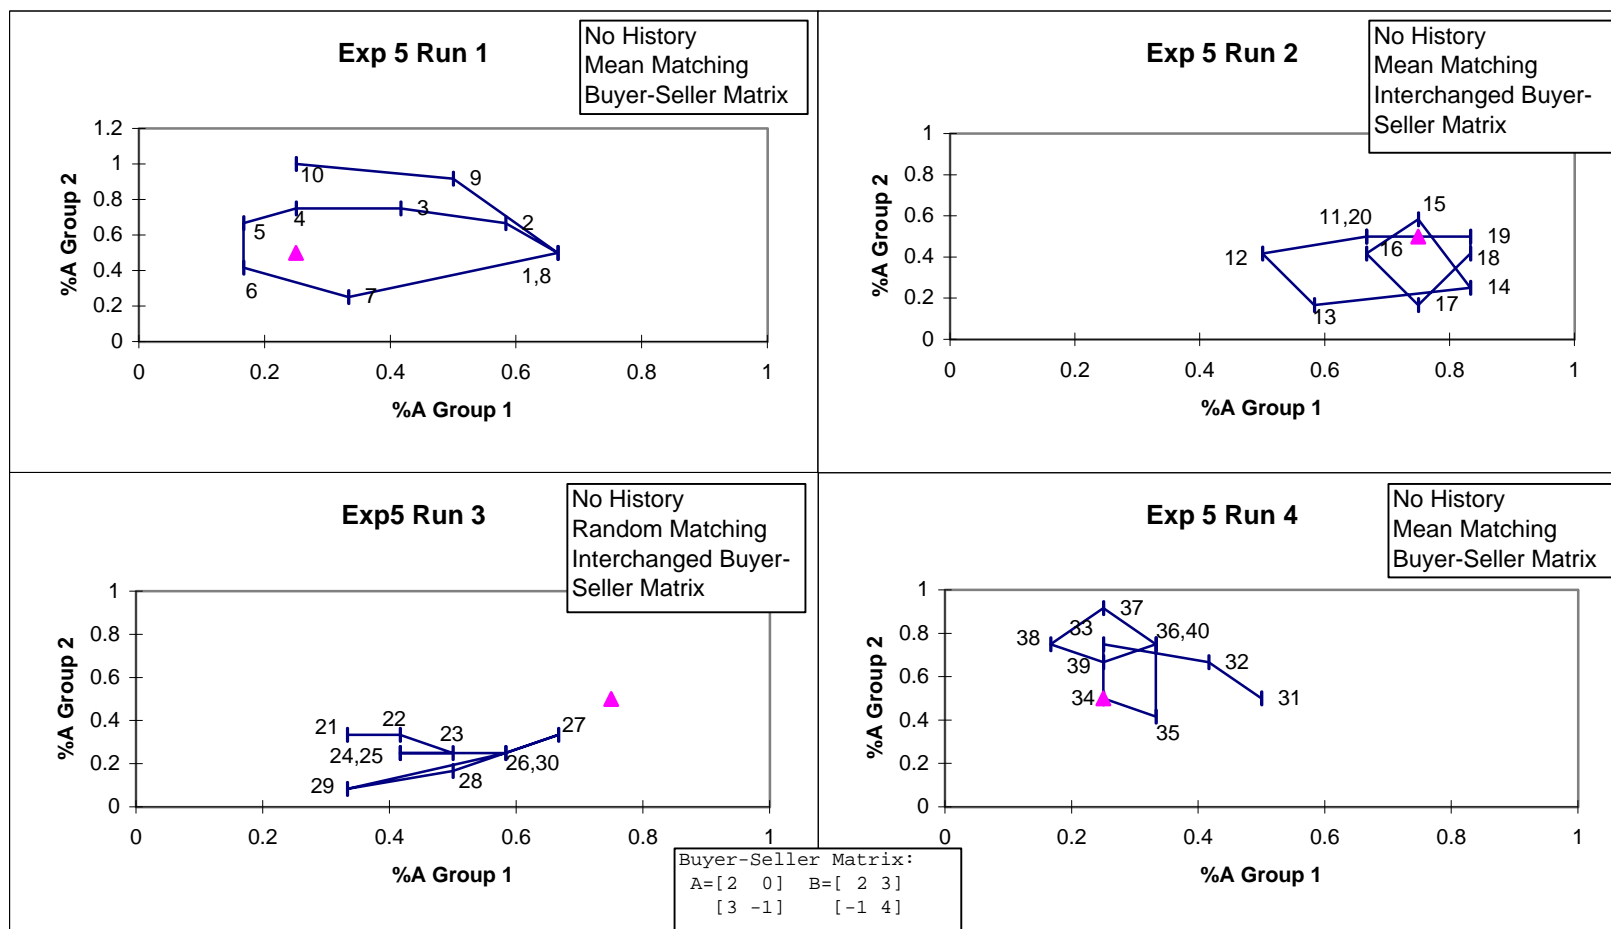


Figure 4: Behavior in the first four runs of exp5, the first 2-population session. All periods use the Buyer-Seller matrix or its interchange, so the unique NE (denoted by \blacktriangle) is at (p,q) \blacktriangle $(.25, .50)$ or, for the interchanged version, at $(.75, .50)$. The graphs show a 2-period moving average of the time path in the unit square. The graph for the first run looks like an unstable counterclockwise spiral diverging from the NE. The second run looks like a tidy counterclockwise double loop around the NE, neither converging nor diverging. The third run uses the random matching protocol; at best there is a weak tendency to drift towards the NE. The fourth run reverts to mean matching and looks like a counterclockwise spiral possibly converging to the NE. In general, mean-matching, (and history) promote faster convergence to a behavioral equilibrium and a closer approximation to NE.

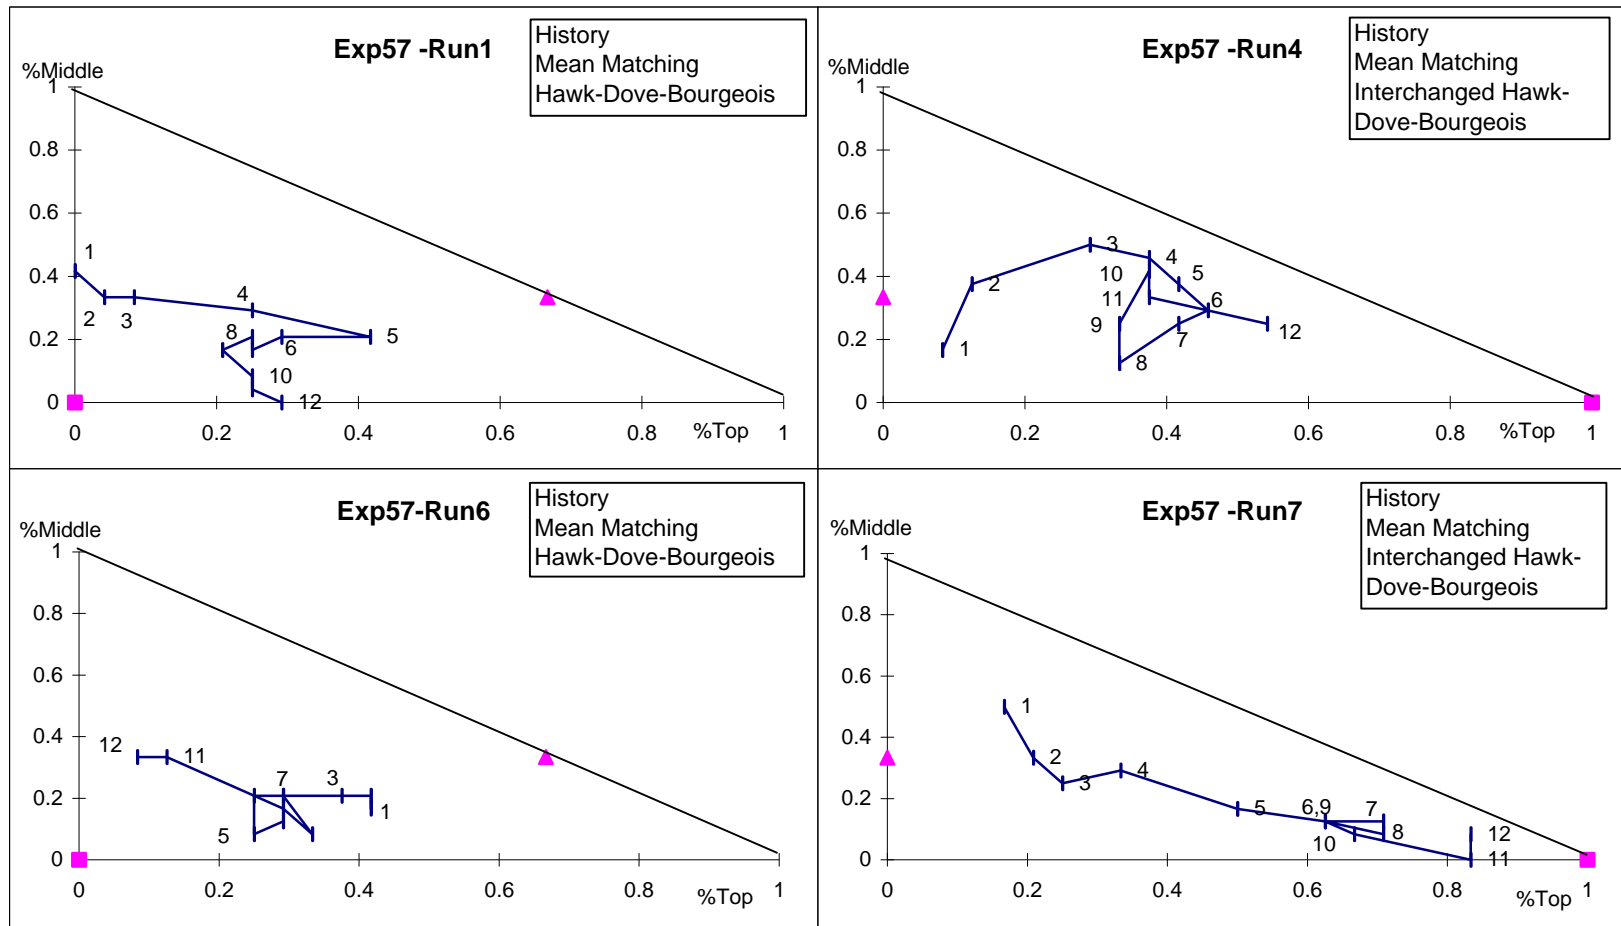


Figure 5: Behavior in the first four HDB runs of exp57. The two NE are at $(p,q) = (0, 0)$ and $(2/3, 1/3)$ (for the interchanged version, at $(1,0)$ and $(0, 1/3)$) and are represented by \blacksquare and \blacktriangle respectively. The graphs show a 2-period moving average. The Run 1 has loose convergence in the second half of the run as does Run 7. Run 4 shows no convergence in either half. Run 6 has loose convergence in both the first and second halves of the experiment. In general, mean-matching (and history) promote faster convergence to a behavioral equilibrium and a closer approximation to NE.