

Separating the Hawks from the Doves: Evidence from Continuous Time Laboratory Games *

Ryan Oprea[†] Keith Henwood[‡] Daniel Friedman[§]

July 5, 2010

Abstract

Human players in our laboratory experiment received flow payoffs over 120 seconds each period from a standard Hawk-Dove bimatrix game played in continuous time. Play converged closely to the symmetric mixed Nash equilibrium under a one-population matching protocol. When the same players were matched in a two-population protocol with the same bimatrix, they showed clear movement towards an asymmetric (and very inequitable) pure Nash equilibrium of the game. These findings support distinctive predictions of evolutionary game theory.

Keywords: Evolutionary dynamics, Hawk-Dove game, game theory, laboratory experiment, continuous time game.

JEL codes: C73, C92, D74

*We are grateful to the National Science Foundation for support under grant SES-0925039, and James Pettit for crucial programming support. We received very useful comments from the audience at the 2009 ESA Meetings at Tucson; seminar audiences at the Cambridge University Business School, the University College London Economics Seminar, and the Stockholm School of Economics; and Juan Carlos Cardenas, Tore Ellingsen, Herb Gintis, Steffen Huck, Chander Velu and Jörgen Weibull. We retain sole responsibility for remaining idiosyncrasies and errors.

[†]Economics Department, University of California, Santa Cruz, CA, 95064. roprea@ucsc.edu

[‡]Economics Department, University of California, Santa Cruz, CA, 95064. keithjhenwood@gmail.com

[§]Economics Department, University of California, Santa Cruz, CA, 95064. dan@ucsc.edu

1 Introduction

Evolutionary game theory offers a principled way to select among multiple Nash equilibria. For example, the famous 2x2 bimatrix game Hawk-Dove (whose variants are sometimes called “Chicken” or “Battle of the Sexes”) has two different asymmetric NE in pure strategies as well as one symmetric NE in mixed strategies. The “mass action” interpretation that Nash first proposed in his dissertation, later elaborated as evolutionary game theory, makes two distinctive predictions: (a) the mixed NE will be selected at the population level when each player interacts with all other players, but (b) one of the pure NE will be selected when row players and column players belong to disjoint populations.

Experimental tests of the theory have so far been sparse, and for good reason. As we will see shortly, standard laboratory procedures involve a limited number of synchronous repetitions of the stage game, while basic evolutionary models involve asynchronous choices in continuous time. Perhaps more importantly, evolutionary game theory provides long run predictions, which may not emerge clearly even in 100 standard repetitions of the stage game.

In this paper we introduce laboratory procedures that mitigate these problems. A new software package called ConG (for **C**ontinuous **G**ames) enables laboratory subjects to make asynchronous decisions in continuous time, to receive instantaneous feedback, and to alter their decisions as often as they like. Behavior can thus settle down within a minute or two, even after covering the gamut of strategic possibilities. Each period in our experiment lasts 120 seconds, allowing considerable stationary repetition (e.g., several runs of 10 periods each) and multiple treatments in a single laboratory session. Thus we can observe reasonable approximations of the “long run” across treatments within a single cohort of subjects.

These procedures enable us to conduct a sharp test of evolutionary selection. Subjects initially play a Hawk-Dove game under either a one-population or two-population matching protocol. Halfway into each session (and without informing subjects or altering the payoff matrix or the feedback display) we switch the matching protocol. This subtle exogenous treatment variation theoretically destabilizes the mixed equilibrium and stabilizes the pure equilibria (or vice versa).

Two other design features are worth highlighting. First, we use a bimatrix with highly in-

equitable payoffs in the asymmetric pure Nash equilibria, giving half of the subjects in the two-population protocol a powerful reason to resist convergence to either predicted equilibrium. Second, by randomly reassigning actions to players at the beginning of each 120 second period, we repeatedly shock the initial conditions, allowing us to address important questions about dynamic adjustment.

As noted earlier, the relevant previous literature is sparse. Van Huyck et al (1995) study a 3x3 bimatrix game with a moderately asymmetric Hawk-Dove-like sub-bimatrix. Using a standard between subjects design with 70 or fewer synchronous repetitions (periods of simultaneous single choices), they find no evidence of convergence with a one-population matching protocol and only weak evidence with a two-population protocol. Friedman (1996) studies a variety of evolutionary games, including two-population Hawk-Dove, each with at most 20 synchronous repetitions. He finds that mean-matching protocols lead to crisper convergence than the standard random pairwise matching. More distantly related literature uses various sorts of evolutionary models to explain persistent departures from Nash equilibrium in coordination games (e.g., Crawford 1991) and from the subgame perfect Nash equilibrium in ultimatum games (e.g., Huck and Oechssler 1999, Guth and Yaari 1992, and Gale, Binmore and Samuelson 1995). A scattering of papers study Hawk-Dove-like games from non-evolutionary perspectives. For example, Neugebauer et al (2008) show that most players do not employ social preferences in their games, and Duffy and Feltovich (forthcoming) examine correlated equilibria. An older literature examined the impact of cheap talk (e.g., Cooper et al. 1989) and forward induction (e.g. Cooper et al. 1993) in such games.

Our presentation begins in the next section with a review of relevant theory. The evolutionary dynamics behind equilibrium selection are often specified via standard replicator systems of ordinary differential equations. More generally, they can be specified via differential inclusions that capture the most basic evolutionary principle, that the fittest (highest payoff) strategy will become more prevalent over time. The theoretical literature on such sign-preserving dynamics is less accessible and so we offer a self-contained exposition of its application to the Hawk-Dove game. The section concludes with a list of the testable predictions arising from both the replicator and sign preserving specifications, and notes how the predictions might go wrong for behavioral or other reasons.

Section 3 describes ConG, the matching protocols, and other aspects of our eight labo-

ratory sessions. The main results are presented in section 4. We observe close convergence to the mixed NE throughout the 10 one-population periods in each session, despite its inefficiency. In 10 two-population periods in the same sessions (with the same subjects and the same payoff bimatrix), we see clear movement towards an asymmetric pure NE, despite its extreme inequity. The replicator and the differential inclusion specifications both help predict which of the two PNE is selected. An examination of the underlying individual behavior discloses a degree of specialization among strategically identical players as well as an interesting interplay between the rates of best response and the incentives to do so.

Following a concluding discussion, Appendix A collects mathematical details and Appendix B shows the instructions to laboratory subjects.

2 Theoretical Considerations

Payoff Bimatrix	NE profile	NE payoffs
$H \begin{pmatrix} 0, 0 & 15, 3 \\ 3, 15 & 9, 9 \end{pmatrix}$	$\frac{2}{3}H + \frac{1}{3}D, \frac{2}{3}H + \frac{1}{3}D$	5,5
	H,D	15,3
	D,H	3,15

Table 1: The Hawk-Dove Game

We use the symmetric payoff bimatrix shown in Table 1. It is an instance of the Maynard Smith (1982) Hawk-Dove game with resource value $v = 12$, cost of conflict $c = 18$ and additive constant $a = 3$. If both players adopt the second strategy (D for Dove), then the resource is divided equally, and the payoff is $a + v/2 = 9$ for each player. If both players adopt the first pure strategy (H for Hawk), then there is a conflict and the (expected) payoff for each is $a + (v - c)/2 = 0$. If they adopt different strategies, then the H player gets the resource and payoff $a + v = 15$ while the D player gets only the additive constant $a = 3$.

As noted in the table, this HD game has two pure asymmetric Nash equilibria, (H,D) and (D,H), and one symmetric mixed NE, $(\frac{2}{3}H + \frac{1}{3}D, \frac{2}{3}H + \frac{1}{3}D)$. Evolutionary game theory selects different equilibria according to whether the matching protocol is one-population or two-population (e.g., Friedman, 1991; Weibull, 1995, pp. 183-186).

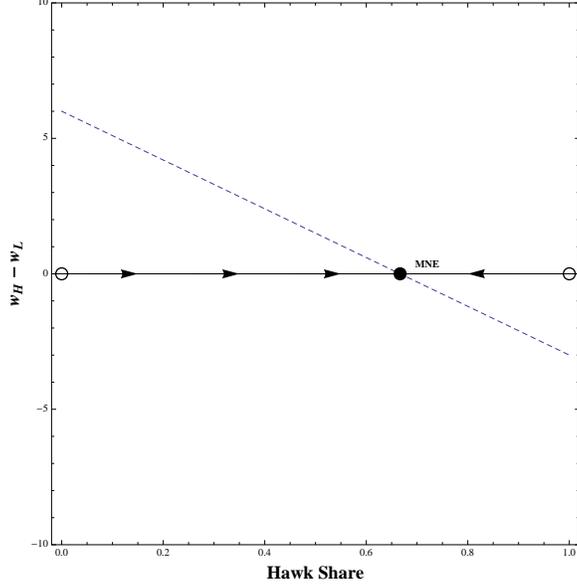


Figure 1: Payoff advantage $\Delta w = w_H - w_D$ for Hawk players in a one-population HD game. Arrows indicate evolutionary dynamics of the Hawk share s .

2.1 One-population dynamics

In the one-population protocol, each player in a “large” population receives the expected payoff when playing a random opponent, or, equivalently, the mean payoff when playing simultaneously the entire population (or a random subset). To spell this out, let s denote the current fraction of Hawk players, so $1 - s$ denotes the fraction of Doves. Then the payoff (or fitness) for a H player is $w_H = 0s + 15(1 - s) = 15 - 15s$ and the payoff for a D player is $w_D = 3s + 9(1 - s) = 9 - 6s$. The payoff difference, the advantage of H over D, is $\Delta w = w_H - w_D = 6 - 9s$.

Solving the interior Nash equilibrium condition $\Delta w = 0$, we obtain the symmetric mixed NE share $s^* = \frac{2}{3}$. The corresponding payoffs are $w_H^* = w_D^* = 5$. The Gini coefficient is zero, since payoffs are equal across all players. Efficiency (NE payoff sum as a fraction of its maximum) is $\frac{5+5}{15+3} = \frac{5}{9}$.

This symmetric mixed NE is stable under any sign preserving dynamic, i.e., any process in which the higher payoff strategy increases its share over time. Figure 1 shows why. Whenever $s < s^*$, we have $\Delta w > 0$ so Hawk has the higher payoff and s increases, while $s > s^*$ implies $\Delta w < 0$ so Dove has the higher payoff and $1 - s$ increases or s decreases.

The best known example of sign preserving dynamics is the replicator of Taylor and Jonker (1978). Replicator dynamics equate the growth rate \dot{s}/s of Hawk play to the Hawk payoff w_H relative to mean payoff $\bar{w} = sw_H + (1-s)w_D$. Thus adjustment is governed by the cubic ordinary differential equation

$$\dot{s} = s(w_H - \bar{w}) = s(1-s)\Delta w(s) = 3s(1-s)(2-3s). \quad (1)$$

Note that the sign of $\Delta w(s)$ determines the sign of the right hand side of (1), so the general result on sign preserving dynamics tells us that the interior steady state $s^* = \frac{2}{3}$ is stable. The other steady states $s = 0, 1$ of (1) are easily seen to be unstable.

2.2 Two-population dynamics

In the two-population protocol, there is a large population of row players (call it population $i = 1$) and a separate large population of column players ($i = 2$). Row players play only column players, and vice-versa.

Let s_i denote the share of Hawks in population $i = 1, 2$. Now the payoff functions are $w_{H1} = 0s_2 + 15(1-s_2)$ and $w_{D1} = 3s_2 + 9(1-s_2)$ in population $i = 1$, so $\Delta w_1 = 6 - 9s_2$. Similarly, in the other population we have $\Delta w_2 = 6 - 9s_1$. Hence replicator dynamics are given by the coupled cubic ODEs

$$\dot{s}_1 = s_1(w_{H1} - \bar{w}_1) = s_1(1-s_1)\Delta w(s_2) = 3s_1(1-s_1)(2-3s_2), \quad (2)$$

$$\dot{s}_2 = s_2(w_{H2} - \bar{w}_2) = s_2(1-s_2)\Delta w(s_1) = 3s_2(1-s_2)(2-3s_1). \quad (3)$$

The phase portrait is shown in Panel A of Figure 2. The interior mixed NE now is a saddle point, hence unstable, while both of the asymmetric pure NE are asymptotically stable. Efficiency at either PNE is $\frac{15+3}{9+9} = 1.0$, but the very unequal payoffs across the two populations of equal mass produce a large Gini coefficient of $1/3$. The other rest points, at $(s_1, s_2) = (0, 0)$ and $(1, 1)$ are unstable. Appendix A verifies the Gini coefficient and uses standard linearization techniques to verify the stability results.

Replicator dynamics are a handy example. They are canonical for haploid population genetics (e.g., Taylor and Jonker, 1978) and can also arise in social adjustment processes based

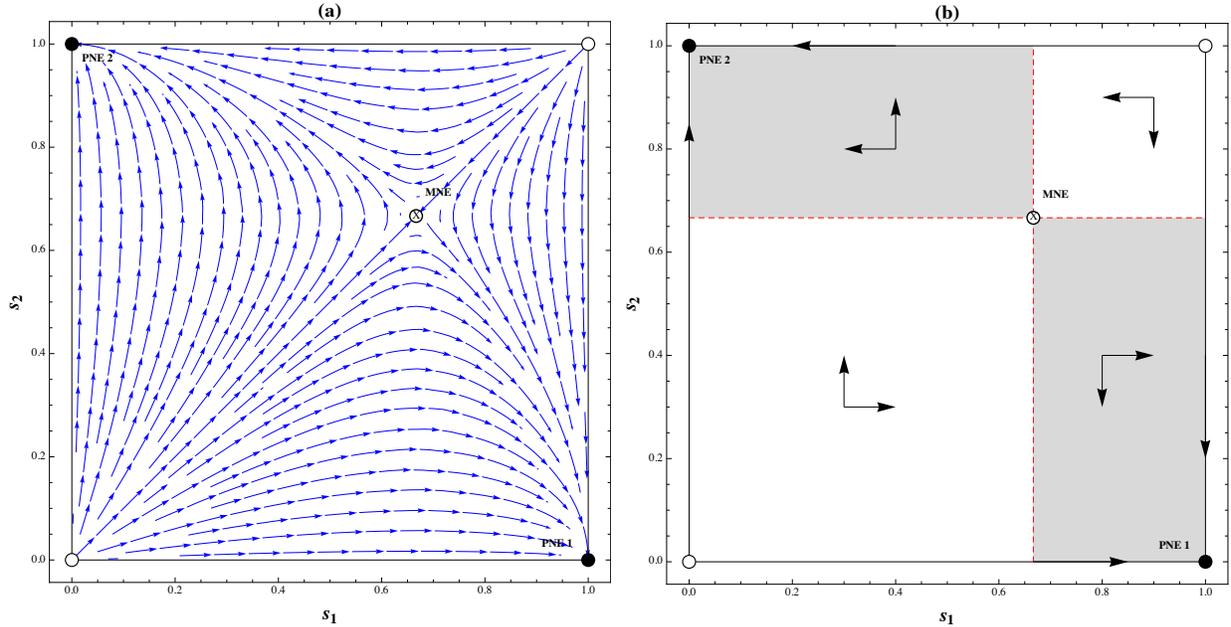


Figure 2: Dynamics for two populations. Phase portrait for replicator dynamics (Panel A), and differential inclusions for sign-preserving dynamics (Panel B). Shaded rectangles are absorbing while unshaded rectangles are transient under sign preserving dynamics.

on imitation (e.g., Schlag, 1998; Björnerstedt and Weibull, 1996). Of course, there are many other plausible adjustment processes for humans that differ markedly from replicator (e.g., Friedman, 1991; Sandholm, 2010, Ch. 5-6), including most learning processes.

The advantage of sign preserving dynamics is that they capture a very broad set of plausible adjustment processes—essentially those for which current material payoffs, as specified here by the payoff bimatrix, dominate other considerations. But do the replicator stability results generalize to this broader class?

The answer is affirmative, subject to a few nuances. To demonstrate, we will use differential inclusions (Aubin and Cellina, 1984; Sandholm, 2010, Ch 6A), previously known as cone fields (Smale, 1976). The rest of the subsection sketches the main ideas and the geometric intuition, while Appendix A.1 includes the technical details.

As with a system of ordinary differential equations, the two-population adjustment process starting from any given initial condition $(s_1(0), s_2(0)) = (s_{1o}, s_{2o})$ is described by a continuous trajectory $\{(s_1(t), s_2(t)) : t \in [0, \infty)\}$. However, for a differential inclusion, the tangent vector (\dot{s}_1, \dot{s}_2) along the trajectory is required only to be included in some vector subspace;

it need not satisfy a particular system of equations. For sign preserving dynamics, that subspace is the orthant defined by the signs of $(\Delta w_1, \Delta w_2)$.

Panel B of Figure 2 helps us visualize the resulting behavior. The horizontal and vertical lines through $(s_1^*, s_2^*) = (\frac{2}{3}, \frac{2}{3})$ divide the state space into four rectangles. At any point in the open Southwestern rectangle (or on its outer edges, on the horizontal and vertical axes), both Δw_i 's are positive, so sign preserving dynamics will produce tangent vectors in the positive orthant \mathfrak{R}_{++} . Hence the fraction of Hawks will increase in both populations, and the trajectory $(s_1(t), s_2(t))$ will move Northeast until it exits this rectangle.

If the state exits across the vertical boundary $s_1 = \frac{2}{3}, s_2 \in [0, \frac{2}{3})$, then it enters the Southeastern rectangle, where $\Delta w_1 > 0 > \Delta w_2$. Here sign preserving dynamics require that (\dot{s}_1, \dot{s}_2) lie in the SE orthant \mathfrak{R}_{+-} , so the trajectory heads southeast and can never leave the Southeastern rectangle. Either (a) it hits the outer edge $s_1 = 1, s_2 \in [0, \frac{2}{3}]$ from which it must converge to the pure NE $(1, 0)$ because $\dot{s}_1 = 0$ and $\dot{s}_2 < 0$ on this edge, or (b) it hits the outer edge $s_1 \in (\frac{2}{3}, 1], s_2 = 0$ from which again it must converge to the pure NE $(1, 0)$, because $\dot{s}_1 > 0 = \dot{s}_2$ on this edge.

On the other hand, if the state exits the SW rectangle across the horizontal boundary $s_1 \in [0, \frac{2}{3}), s_2 = \frac{2}{3}$, then it enters the Northwestern rectangle. Since $\Delta w_1 < 0 < \Delta w_2$ in this rectangle, the adjustment orthant is \mathfrak{R}_{-+} and the trajectory heads northwest. As before, the trajectory must hit the outer boundary and then converge to the corner—in this case, to the pure NE $(0, 1)$.

Trajectories beginning in the NE rectangle are similar to those beginning in the SW. They must eventually enter the SE or NW rectangles and then converge to one of the pure asymmetric NE. If a trajectory happens to hit the interior Nash equilibrium point $(s_1^*, s_2^*) = (\frac{2}{3}, \frac{2}{3})$, then the adjustment direction is in $\mathfrak{R}_{oo} = \{(0, 0)\}$, so it is a steady state. This equilibrium point is dynamically unstable because any of its small open neighborhoods contains points in the NW and in the SE rectangles and, as we have seen, trajectories from such points move away from the interior equilibrium and towards a corner equilibrium.

To summarize, sign preserving dynamics for the two-population protocol of the HD game have two stable equilibria, at the corners of the state space at which one population plays all Hawk and the other population plays all Dove. The basin of attraction (the set of initial

states whose trajectories converge to that equilibrium) for each of these asymmetric pure equilibria includes the neighboring rectangle, and that rectangle is *absorbing* in the sense that no trajectories starting in it can leave it. The other two rectangles (actually, squares in the SW and NE) are *transient* in that all trajectories starting in them eventually exit.

The geometry of convergence is slightly different for the replicator dynamic. Its basins of attraction are the triangles separated by the saddle path $s_1 = s_2$. (It may be worth noting that the one-population protocol can be regarded as restriction to this saddle path and, subject to this restriction, the symmetric interior equilibrium is stable in any sign preserving dynamic.) Each of these basins contains one of the absorbing rectangles plus half of each transient rectangle.

2.3 Testable Predictions

Economic applications of evolutionary games deal with finite populations of human players. Here sampling error and behavioral noise will produce stochastic terms (see Appendix A.2), so the large population, deterministic theory just reviewed should be used to predict central tendencies, not exact behavior.

With this caveat in mind, we obtain the following testable predictions for the Hawk-Dove game played over time in finite human populations.

Prediction 1 . *Under a one-population matching protocol, the average fraction of Hawk play will converge to $s^* = \frac{2}{3}$, while in a two-population protocol with equal population sizes, the overall average fraction will converge to $\bar{s} = \frac{1}{2}$.*

The last part of this prediction is based on the more specific

Prediction 2 . *Under a two-population matching protocol, the average fraction of Hawk play will approach $s_i = 1$ in one population and $s_{-i} = 0$ in the other.*

A feature of Hawk-Dove games is that D-D and H-D pairings efficiently split the available surplus $v > 0$, while H-H pairings incur an efficiency loss c . The asymmetric pure NE therefore are more efficient than the mixed interior equilibrium (1 vs. $1 - \frac{2^2}{3} = \frac{5}{9}$), albeit far less equitable (15:3 vs. 5:5). Hence we have

Prediction 3 . *A two-population matching protocol will lead to greater efficiency and a substantially less equitable distribution of earnings than the one-population protocol.*

Finally, the replicator basins of attraction are triangular and separated by the diagonal $s_1 = s_2$, while convergence in more general sign-preserving dynamics is governed by the absorbing rectangles. These alternative approaches lead respectively to versions (a) and (b) of

Prediction 4 . *Assume $s_{1o} \neq s_{2o}$. The population i that more nearly converges to all-H in a two-population matching protocol is more likely to be the one that (a) has the larger initial fraction s_{io} of H play, or if possible, (b) satisfies the initial condition $s_{io} > \frac{2}{3} > s_{-io}$.*

All of the predictions could fail in the laboratory if the human subjects respond more strongly to other considerations than to personal material payoff. Idiosyncratic tastes or systematic errors are always possible. A smaller-than-predicted fraction of H play could be attributed to risk aversion, since the bimatrix specifies a larger payoff spread for H (0 or 15) than for D (3 or 9). A more interesting possibility is that human subjects may shun the efficient but very inequitable equilibrium payoffs of (3, 15), and instead seek the efficient egalitarian D-D outcome of (9, 9). In this case, the observed fraction of H play would again be smaller than predicted in both matching protocols, and efficiency would be higher than predicted in the one-population protocol.

On the other hand, we could observe a larger fraction of Hawk play than predicted if subjects engage in negative reciprocity. In particular, the two-population game could be seen as a team contest with the winners gaining a prize five times larger than the losers' prize. The result could be an ongoing war of attrition, or vendetta, with very high and persistent levels of H play in both populations.

3 Experimental Design

To test these predictions, we conducted eight continuous time laboratory sessions with 12 (in one session 10) subjects in each. In each of 20 two minute periods in each session, subjects were independently and randomly assigned an initial action and then could change their actions at any time and as often as they liked. They earned flow payoffs generated by the

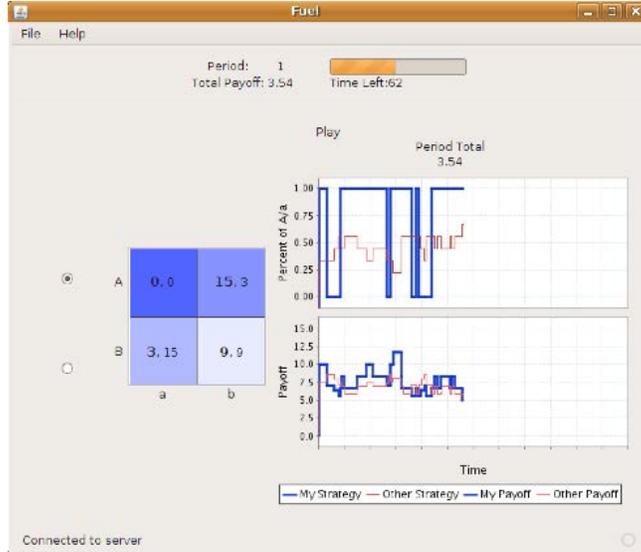


Figure 3: Screenshot of continuous time display.

matrix in Table 1.

This continuous time implementation was enabled by a new software package called ConG.¹ A screenshot of the user interface is provided in Figure 3. On the left side of the screen is a reproduction of the payoff matrix; the row corresponding to the subject's current action is highlighted in blue (every subject is framed as the row player). Moreover, blue shading on the columns becomes darker as more subjects from the opposing population play the corresponding action.

On the right side of the screen are two time series evolving over the course of the period, the top representing actions and the bottom payoffs. In each case the subject's own action is represented in blue while the mean of the opposing population is represented in red. The interface also displays the time remaining in the period, period earnings, and cumulative earnings.

Subjects were assigned (without their knowledge) either to population $i = 1$ or to population $i = 2$ at the beginning of the session. We study two treatments, varied within session. In the one population treatment, denoted 1P, each player j is matched each period with each of the $n - 1$ other players. Her instantaneous payoff when n_{-jH} of the other players are choosing A (or Hawk) is $15 - 15 \frac{n_{-jH}}{n-1}$ if she plays A and is $9 - 6 \frac{n_{-jH}}{n-1}$ if she plays B. In the two

¹Friedman and Oprea (2010) use ConG in a very different setting. They study pairwise play of prisoner's dilemmas in continuous time and find very high rates of cooperation.

population treatment, denoted 2P, each subject j in population $i(j)$ is matched only with the $n/2$ subjects in the other population $-i$ and earns instantaneous payoff $15 - 15\frac{n_{-iH}}{n/2}$ for A or $9 - 6\frac{n_{-iH}}{n/2}$ for B, where n_{-iH} is the current number of Hawks in the other population $-i$.

Instantaneous payoffs are accumulated over each two minute period. For example, if all players independently spend half their time as Hawks (A), then each player earns $\frac{1}{2}(15 - 15 \cdot \frac{1}{2}) + \frac{1}{2}(9 - 6 \cdot \frac{1}{2}) = 6.75$ points per period.

At the beginning of each period subjects were randomly assigned an initial action. We will exploit this exogenous variation to study how initial conditions impact convergence.

Instructions mentioned the two matching protocols, but subjects were never told which was in force, and the interface offered no clues. Subjects observed only mean payoffs and mean action choices of their counterpart population and because initial actions were shocked to a random position at the beginning of each period it was impossible to make firm conclusions about structural changes at the beginning of period 11. This design choice further strengthens our test. Because subjects had no opportunity to consciously coordinate on a change of strategy with the change of treatment, treatment effects can only emerge by the sorts of adaptive forces specified by the theory.

In sessions 1-4 subjects interacted under treatment 1P for the first 10 periods and 2P for the second half. In sessions 5-8 this ordering was reversed.²

All sessions were conducted using inexperienced undergraduate subjects at the University of California, Santa Cruz. Instructions, read aloud to subjects, are reproduced in Appendix C. Sessions lasted roughly 75 minutes and earnings, including a five dollar show up fee, averaged \$17.

4 Results

We begin with an overview and tests of the first three predictions. Later subsections examine the fourth prediction and the underlying individual behavior. For tractability, the reported

²Supplementary data analysis, available on request, confirms that the order of treatments has no qualitative effect on our main results.

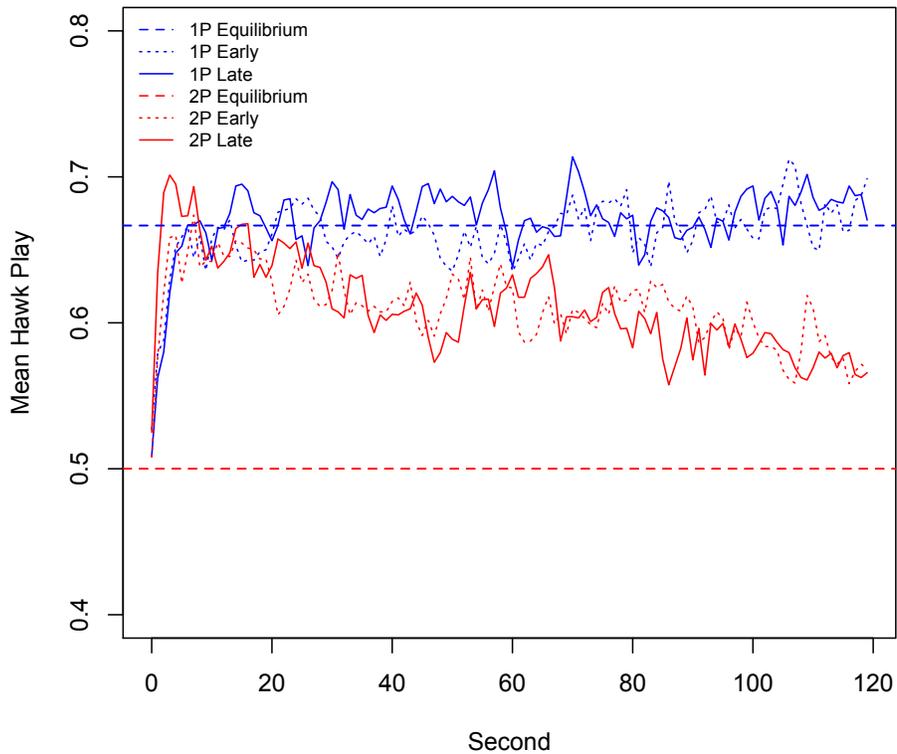


Figure 4: Mean rates of Hawk play within period.

data are samples from each one-second interval.

4.1 Within Period Behavior

Figure 4 plots, second by second, the mean fraction of Hawk play over all 160 periods of data. Blue (red) lines denote 1P (2P) periods, and dotted (solid) lines are from the first five (last five) periods in the treatment. Due to the random assignment of initial actions, all lines start near 0.5, and all rise quickly. The 1P lines reach the predicted MNE value of 0.67 within a few seconds and remain in that vicinity for the rest of the period. The 2P lines also jump quickly to 0.67 or above, but then trend downward over the rest of the period, eventually covering about half the distance to PNE level of 0.5.

	(1)	(2)	(3)	(4)
	Hawk	Separation	Efficiency	Gini
Intercept	65.023*** (0.900)	21.701*** (0.871)	59.305*** (1.818)	4.45*** (0.295)
$P2$	0.453 (0.301)	11.350*** (1.864)	2.421 (1.923)	16.76*** (1.489)
t	0.027*** (0.003)	0.013 (0.008)	-0.031*** (0.007)	-0.007*** (0.001)
$P2 \times t$	-0.086*** (0.004)	0.214*** (0.035)	0.165*** (0.032)	0.06*** (0.005)
At t=120 seconds				
1P	69.06***	23.26***	54.62***	3.67***
2P	57.79***	60.27***	76.80***	27.51***

Table 2: Regression results for behavior within period. In all cases, standard errors are robust and clustered at the session level. Estimates and standard errors are multiplied by 100 for legibility and can be interpreted as percentages.

To obtain more quantitative evidence, we estimate the following equation:

$$s_{Hjkt} = \beta_0 + \beta_1 P2_{jk} + \beta_2 t + \beta_3 P2_{jk} \times t + \epsilon_{jkt} \quad (4)$$

where s_{Hjkt} is the observed fraction of Hawk play in second $t \in \{1, 2, \dots, 120\}$ of period k of session j , while $P2$ is an indicator for the two-population treatment and ϵ is assumed to be Gaussian distributed. Standard errors are clustered at the session level to control for within-group dependence.

Column (1) of Table 2 collects the coefficient estimates. The entry for t indicates a slight positive trend in the 1P treatment, but the interaction entry indicates that the trend is reversed and much stronger in 2P. The last two lines show the estimated rates of Hawk play at the end of the period. A Wald test ($p = 0.000$) confirms that these final rates differ significantly across treatments.

Thus the evidence supports Prediction 1 and can be summarized as follows.

Result 1 *Average rates of Hawk play quickly and closely approximate the mixed NE value of 2/3 in the 1P treatment, and they move decisively towards pure NE value of 1/2 in the 2P treatment.*

Recall that Prediction 2 is that, in treatment 2P, the rates of H play in the two populations

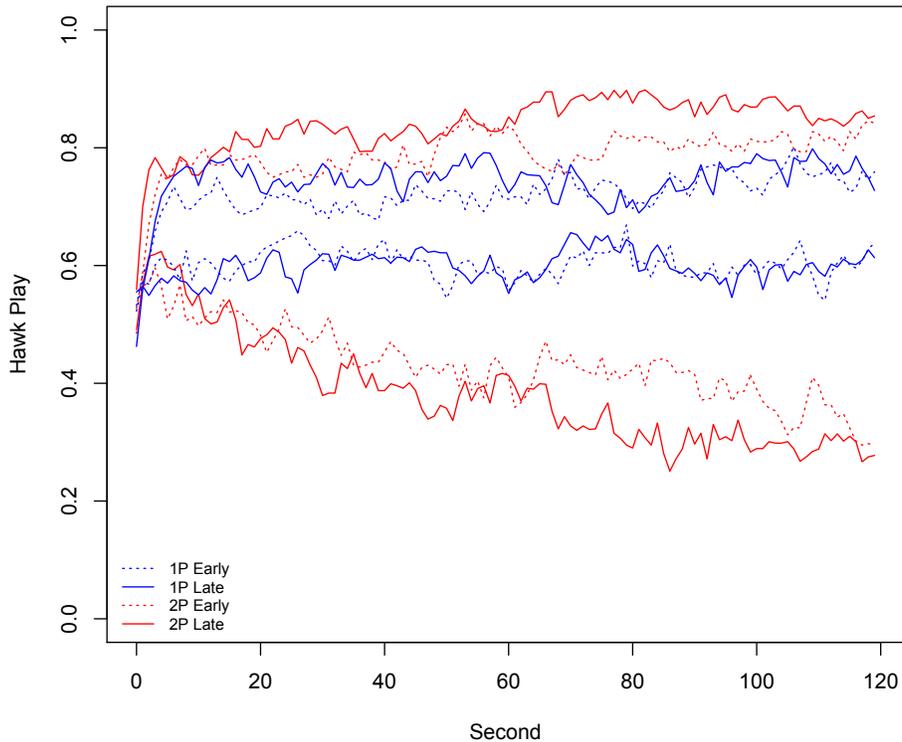


Figure 5: Mean rates of Hawk play by second within period. Separate series are plotted for the (weakly) more Hawkish and less Hawkish group to visualize the degree of separation in rates across groups.

will diverge towards 0 and 1. Of course, in 1P, the rates should not differ systematically because group identity is meaningless in that matching protocol. However, Appendix A notes that, given binomial error in the 1P mixed NE, the group with the higher rate should have on average 0.21 more Hawk play.

Figure 5 gives an overview of the relevant evidence. Using the same coloring and line weighting schemes as before, it displays separate lines for the group with the higher rate and the group with the lower rate of Hawk play that period. The lines in the 1P (blue) treatment bracket the mixed NE value of $\frac{2}{3}$ and remain separated by almost exactly the predicted 0.21.

By contrast, the 2P (red) lines diverge as the Hawkish group on average becomes increasingly Hawkish and the Dovish group gradually becomes more Dovish. This predicted separation

is not complete; on average between 1 and 2 players deviate from pure NE in the Dovish population and 1 player deviates in the Hawkish population. Nonetheless, the increasing separation across groups confirms Prediction 2. It is particularly striking in view of the fact that group assignments are constant across the two treatments— the same groups of individual subjects make similar choices in 1P but make divergent choices in 2P, no matter which treatment comes first.

More quantitative evidence comes from a regression similar to (4) above. Now the dependent variable is separation $\Delta s_{jkt} = s_{[1](jk)}(t) - s_{[2](jk)}(t) \geq 0$, the difference in second t between the observed fraction of Hawk play in the group $[1](jk)$ with the larger average in period k of session j and that observed in the other group $[2](jk)$.

$$\Delta s_{jkt} = \beta_0 + \beta_1 P2_{jk} + \beta_2 t + \beta_3 P2_{jk} \times t + \epsilon_{jkt} \quad (5)$$

Column (2) of Table 2 reports the coefficient estimates. The intercept measures the initial separation in 1P, and its value of 0.21 is precisely as predicted in a mixed NE. The small and insignificant t coefficient indicates no systematic separation trend in 1P, but the large positive interaction coefficient indicates a very strong trend in 2P. By the end of the period in 1P, on average there is about one extra H player (23 percent of the population of 5 or 6 players) in the more Hawkish population, while in 2P the average final difference is at least three players (60 percent). The difference across treatments is significant at the one percent level. The evidence thus provides support for Prediction 2, and can be summarized as follows.

Result 2 *The difference in Hawk play between the two groups is small and trendless in 1P, while in 2P it increases steadily and becomes quite large and significant.*

Prediction 3 concerns the impact on absolute and relative welfare. Recall that efficiency (Eff = total realized payoff as a fraction of the maximum possible) is predicted to be $\frac{5}{9}$ in 1P and to be 1.0 in 2P. Panel (a) of Figure 6 shows rates of efficiency on average at each second for both treatments. Efficiency starts at 75% (because the H-H encounter rate begins at 25% due to the random assignment of initial actions) but in 1P periods it drops almost immediately to the predicted level and stays there. It shows a similar immediate drop in 2P periods, but then moves decisively higher and approaches 80 percent by the end of later

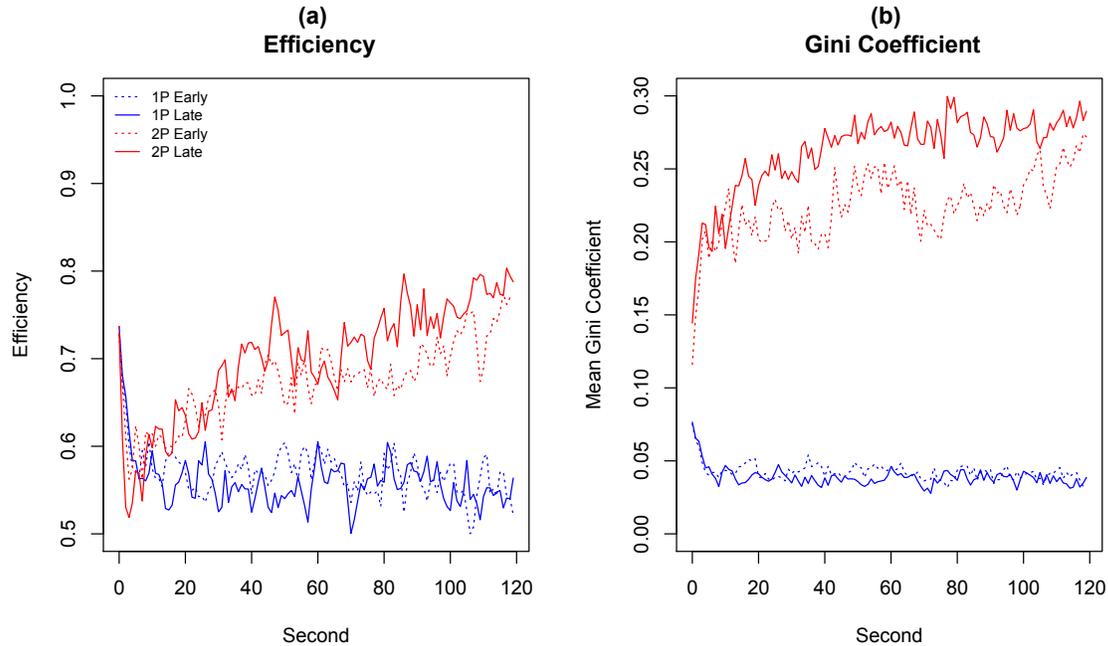


Figure 6: Mean efficiency (Panel a) and inequality (Panel b) within period.

periods (solid red line).

The usual regression, with efficiency as the dependent variable,

$$\text{Eff}_{jkt} = \beta_0 + \beta_1 P2_{jk} + \beta_2 t + \beta_3 P2_{jk} \times t + \epsilon_{jkt}, \quad (6)$$

yields the coefficient estimates shown in column (3) of Table 2. The small but significantly negative t coefficient probably reflects the impact of the strategy initialization; the final estimate of 55% is right on target for P1. Average efficiency in 2P reaches a final value of 77% and is trending strongly upward.

These efficiency gains come at a substantial cost to equity. Panel (b) of Figure 6 reports the Gini coefficient, the standard inequality measure. As predicted, it falls quickly to a very low level in 1P, but increases immediately in 2P and continues to rise modestly over time and across periods, eventually reaching 27.5%, not far below the predicted level of 33%.

These impressions are tested in the regression

$$\text{Gini}_{jkt} = \beta_0 + \beta_1 P2_{jk} + \beta_2 t + \beta_3 P2_{jk} \times t + \epsilon_{jkt}, \quad (7)$$

with results shown in the last column of Table 2. It confirms that the differences between the 1P and 2P treatments are very large, increasing, significant, and in the predicted direction.

Result 3 *The two-population treatment generates substantially greater average inequity and substantially greater average efficiency than the one-population treatment, and these differences increase over the course of the period.*

4.2 Initial Conditions and Equilibrium Selection

Prediction 4 says that pure Nash equilibrium PNE i is selected in a given 2P period if the initial conditions (s_{1o}, s_{2o}) lie (a) on same side of the diagonal $s_{1o} = s_{2o}$ as PNE i , or at least (b) in the absorbing rectangle containing PNE i .

The prediction concerns the final state, and as its empirical counterpart we chose (s_{1Fjk}, s_{2Fjk}) = average play over the last 10 seconds of period k of session j . As a robustness check, we also looked at the last 20 seconds and found no substantive difference.

The prediction is contingent on the initial state, and there are two obvious choices for its empirical counterpart: (i) the initial state determined by the realized randomization that period, or (ii) the final state reached in the previous period.

The relevant data come from sessions $j = 1, \dots, 8$, each of which has 9 periods in treatment 2P with a previous 2P period. In 71 of these 72 periods, the final state (s_{1Fjk}, s_{2Fjk}) lies inside one of the absorbing rectangles, and the exceptional period is a near-miss, so there is no ambiguity in deciding whether PNE1 or PNE2 is being selected.

In testing the prediction, we focus on the final difference in Hawk shares, $D_{jk} = s_{1Fjk} - s_{2Fjk} \in [-1, 1]$. The value $D_{jk} = 1$ indicates full convergence to PNE1 (all Hawk in population 1) in period k of session j , while $D_{jk} = -1$ indicates full convergence to PNE2 (all Hawk in population 2). Intermediate positive (negative) values indicate incomplete convergence to PNE1 (to PNE2). We implement (ii) above as last period's final difference D_{jk-1} , and implement (i) as the random initialization difference $D_{ojk} = s_{1ojk} - s_{2ojk} \in [-1, 1]$.

For the sake of completeness, we estimate the equation

$$D_{jk} = \beta_0 + \beta_1 P2 + \beta_2 D_{jk-1} + \beta_3 D_{jk-1} \times P2 + \beta_4 D_{ojk} + \beta_5 D_{ojk} \times P2 + \epsilon_{jk} \quad (8)$$

Variable	Estimate	Standard Error
Intercept	-0.006	0.050
$P2$	0.073	0.059
D_o	-0.077	0.078
$P2 \times D_o$	0.436**	0.167
D_{k-1}	0.204	0.156
$P2 \times D_{k-1}$	0.395*	0.215

Table 3: Coefficient estimates for equation (8). Standard errors are clustered at the session level. One, two and three asterisks indicate significance at the 10, 5 and 1 percent levels

on the 72 data points in 2P together with the 72 corresponding data points in 1P.

Table 3 collects the results. The intercept and linear terms have insignificant coefficients, consistent with the evolutionary game theory implication that initial conditions have no lasting impact in the 1P treatment. The impact of initial conditions in 2P is measured by the sums $D_o + P2 \times D_o = 0.359^{**}$ and $D_{k-1} + P2 \times D_{k-1} = 0.598^{***}$. Both impacts are positive and significant, especially that of (ii), the Hawk difference last period.

Our interpretation is that the random initialization does have some influence on subsequent behavior in 2P, but that behavior responds somewhat more strongly to the shared experience of what happened last period. If Hawkishness in group 1 relative to group 2 increased by some amount x this period then, other things equal, we'd expect about $0.60x$ more relative Hawkishness next period. The same increase x in random initial assignments would push final behavior in the same direction but only by about $0.36x$.

The variable D_{jk-1} captures groups' lagged behavior and it may partly reflect the fact that some groups simply have more aggressive players than others. The insignificance of this variable in 1P suggests that such group composition is not the dominant force even in 2P.³

Thus the data are generally consistent with Prediction 4a, as summarized in

Result 4 . *The final state in the previous period has a modest and statistically insignificant impact on the current final state in 1P but has a substantial and positive effect in 2P. The same is true (to a slightly lesser extent) of the random initialization.*

³Indeed, as a robustness check (available upon request), we estimated a version of (8) using only 2P data but with the mean value of D from corresponding 1P periods included as an explanatory variable to control for group differences in intrinsic Hawkishness. Even with this control, D_{jk-1} remains large and significant, again suggesting that composition effects are not the main driver of the estimates in Table 3.

	Different	Same	Total
Absorbing	11/20=0.55	17/19=0.89	28/39=0.72
Transitory	1/11=0.09	6/8=0.75	7/19=0.37
Total	12/31=0.39	23/27=0.85	35/58=0.60

Table 4: Contingency Table. Entries show the number of correct predictions of the final state’s absorbing rectangle as a fraction of the number of periods in each contingency combination. The contingencies are whether the random initialization is in an an Absorbing or Transitory rectangle, and whether it is on the Same or Different side of the diagonal as last period’s final state.

Is it really the distance from the diagonal that matters in 2P (as suggested by the regression and by replicator dynamics) or is it just the absorbing rectangle (as in Prediction 4b and sign preserving dynamics)? Variable D_{jk-1} can’t tell us, because the (last period) final state virtually always belongs to the absorbing rectangle. So the empirical question becomes: how much more predictive is the random initialization when it lies in an absorbing rectangle?

A contingency table provides some direct evidence. Of the 72 observations in 2P, we drop the 14 in which the initialization is on the diagonal and thus uninformative. For the remaining 58 observations, Table 4 shows that, when it disagrees with the previous final state, the current random initialization correctly predicts the selected equilibrium in about half the observations (11 of 20) when it lies in the absorbing rectangle, but only in 1 of 11 when it does not. When it agrees with the previous final state, the current random initialization correctly predicts the selected equilibrium in all but four observations (23 of 27); here it makes little difference whether or not it lies in an absorbing rectangle.

Our interpretation is that the initialization has no influence unless it is in the absorbing rectangle and disagrees with the previous final state. In other cases, the previous final state (which in each of the 58 observations lies in an absorbing rectangle) seems the best predictor.⁴ Hence the evidence seems entirely consistent with Prediction 4b, and may be encapsulated as follows.

Result 5 *As in sign preserving dynamics, the initial condition has predictive power when it lies in an absorbing rectangle. We found no evidence of predictive power for initial conditions that lie in a transitory rectangle.*

⁴A dummy variable regression (available on request) confirms that these observations are statistically significant.

4.3 Individual Behavior and Best Response

The theory of evolutionary games is concerned with population distributions and has little to say about individual behavior. Of course, it is individual decisions that create the population distributions observed in our experiment, and it is of some independent interest to explore their nature.

For example, Prediction 1 says, correctly, that in a one-population treatment, the overall fraction of Hawk play should stabilize at the symmetric MNE value of $\frac{2}{3}$, but it doesn't say how. Do all individuals try to employ the same mixed strategy $\frac{2}{3}H + \frac{1}{3}D$, say by choosing independently from the mix at regular intervals? Or by following a Markov switching process with a faster transition rate from D than from H? Or do different individuals specialize differently, with about twice as many consistently playing H as consistently playing D?⁵

Let us say that a player is a *strong specialist* if she spends at least 90% of the period in one of the strategies, and a *specialist* if the fraction is at least 80%. Figure 7 shows that half the Hawkish players in 2P are strong specialists and about 2/3 are specialists. Even in 1P, half the Hawkish players are specialists and about 1/3 are strong specialists. On the other hand, less than 1/5 of Dovish players are strong specialists in either treatment, and the median rate of D play among Dovish players is under 65% in 1P and just a bit over 70% in 2P. Since there are more Hawkish than Dovish players, we conclude that overall about half the players are specialists.

It is reasonable to suppose that specialists seldom switch strategies, but how about non-specialists? Panel (a) of Figure 8 shows that on average, about 10 percent of subjects switch each second throughout the period (after an initial flurry) in both population treatments and in both early and late periods. Panel (b) shows falling switch rates in 2P—the median player switches 10 times in the first period, and only 5 times by the last—and a more modest decline in switch rates in 1P periods. Panel (c) shows the other side of the same coin, the median proportion of time a subject spends in their more favored action. This rises modestly over time and, as we have already seen, is a bit higher in 2P. Thus non-specialists switch fairly often, even in later periods when they become less common.

⁵Friedman (1996) argues that his subjects in discrete time games mostly do specialize, and that near interior NE this can be understood as a variant Harsanyi's (1973) notion of purification.

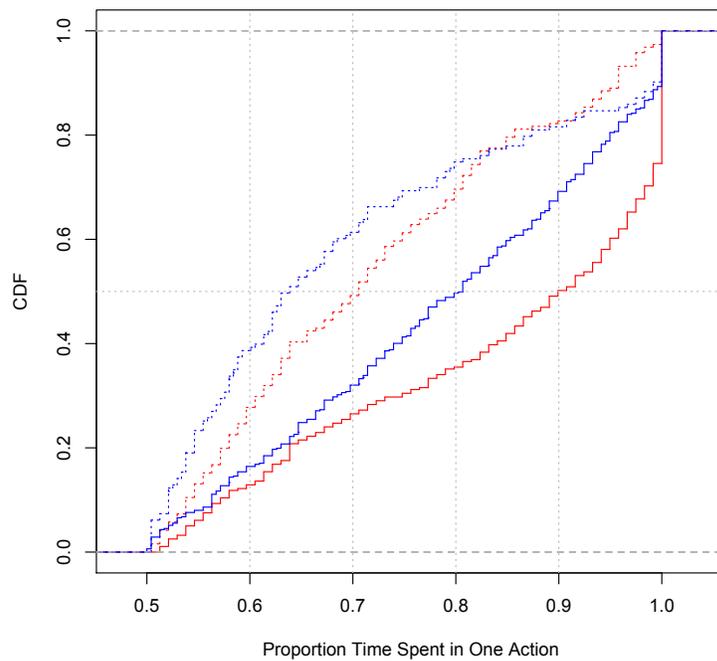


Figure 7: Cumulative distribution functions of specialization rates over last 5 periods. The dotted lines (blue for 1P treatment, red for 2P) show the fraction of time in Dove for subjects who played Dove more than half the time, while solid lines show the fraction of time in Hawk by Hawkish players.

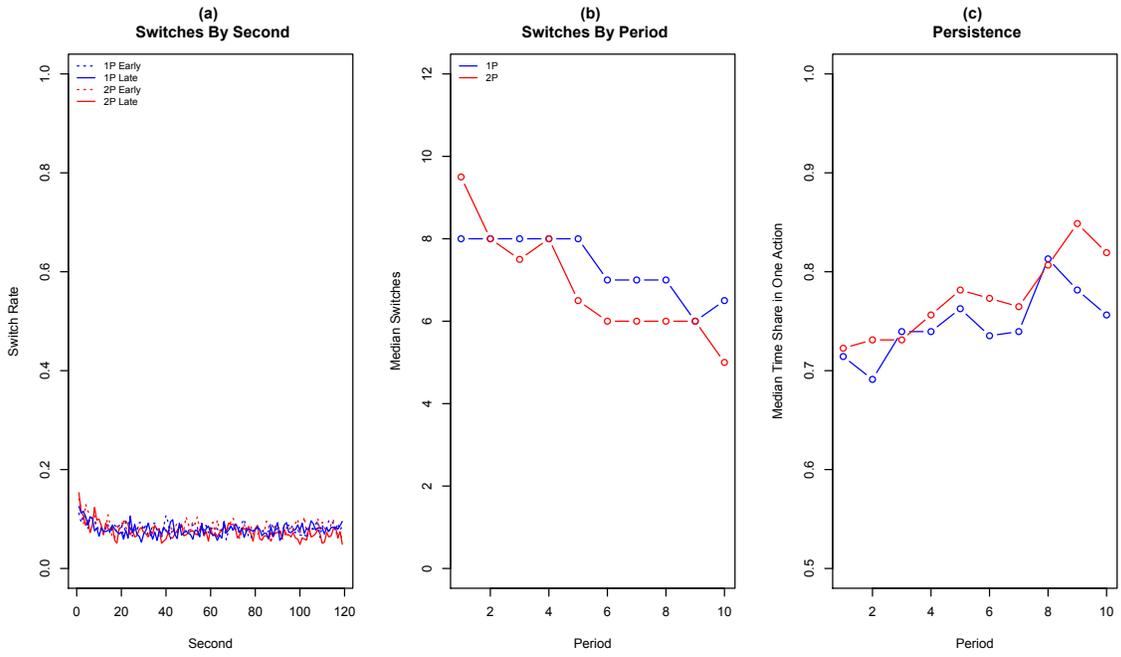


Figure 8: Switch rates within (Panel a) and across periods (Panels b and c). Blue (red) indicates 1P periods (2P periods).

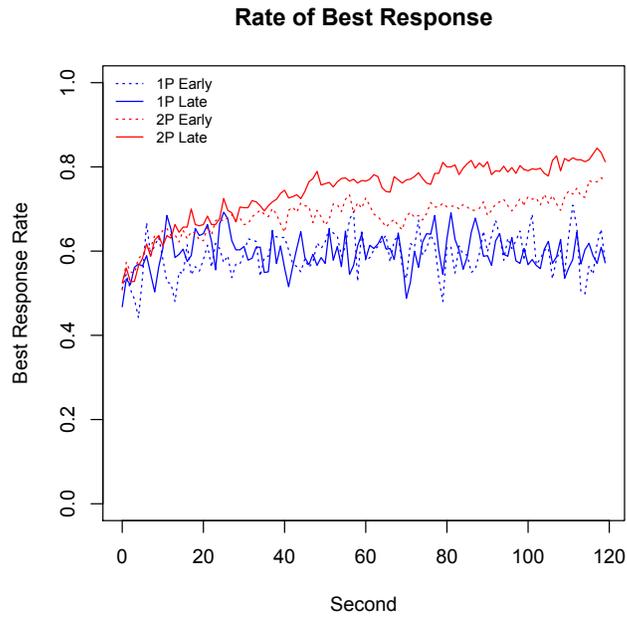


Figure 9: Best response rates within period.

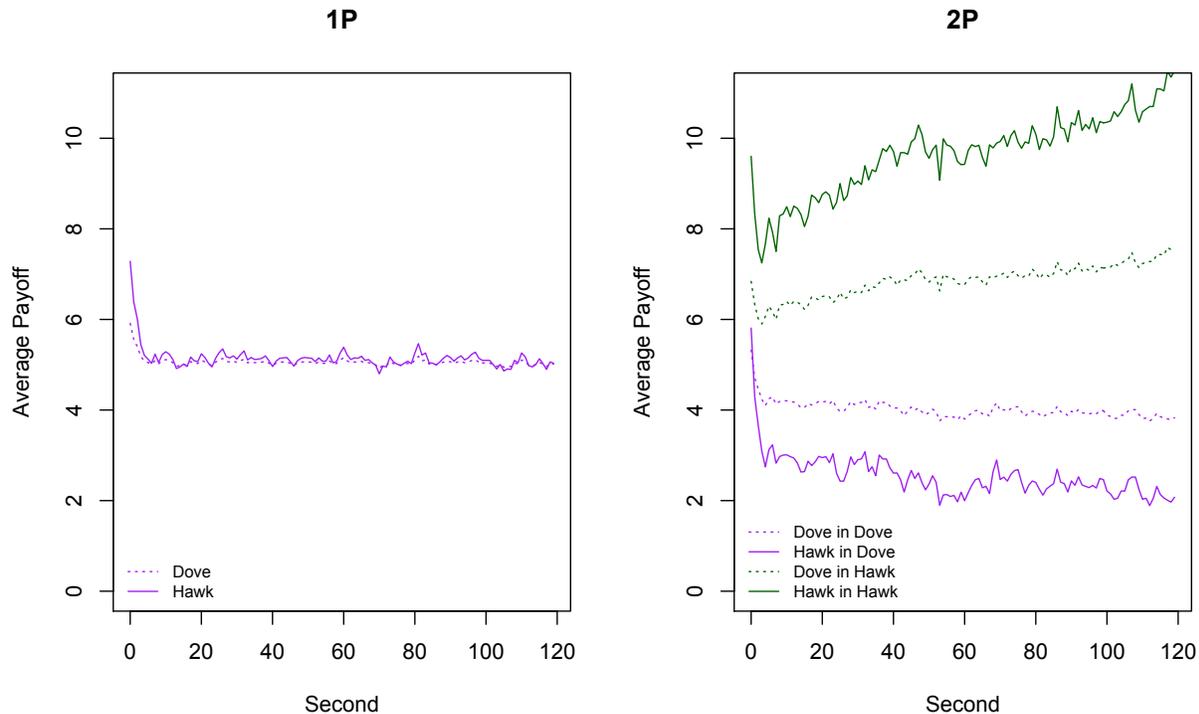


Figure 10: Payoffs by treatment, action and population tendency plotted at each second.

The main conclusions on individual behavior may be summarized as follows.

Result 6 *Individual player behavior is inconsistent with symmetric stationary mixing. Switching becomes less common in later periods periods, and eventually a majority of players specialize in Dove or (more commonly) Hawk.*

Underlying evolutionary dynamics, as we saw earlier, is the principle that players tend to switch to the higher payoff strategy over time. But there is a countervailing force, at least in treatment 2P—human players may play H even when it has a lower payoff in order to hinder the other population and perhaps help their own. In retrospect, some evidence for this “vendetta hypothesis” can be found in Figure 4: the initial rise in H play is steepest and highest in later 2P periods.

Figure 9 shows when players actually employ the higher payoff strategy. Perhaps surprisingly, it shows that only about 60% of players do so in 1P after the first few seconds, and that there are no apparent trends within or across periods. In 2P, the best response rate also begins around 50% (as it must, due to random initialization), but it trends upward both

within period and across periods, and eventually reaches about 80%.

Figure 10 helps resolve the puzzle. It shows that in 1P the two strategies give essentially the same payoff after the first few seconds, so there is hardly any advantage to best responding. In 2P, however, the Hawkish population has a clear incentive to best respond, and it increases over time. This surely helps explain the increase in Hawkish play in this population. The advantage to playing Dove in the Dovish population is smaller and apparently trendless. This provides a self-interest explanation of why we see stronger convergence in the Hawk population than Dove population.

To summarize,

Result 7 *In the one-population treatment, players on average are nearly indifferent between actions and the best response rate is static and not large. In the two-population treatment, the incentive to best respond, and the rates of best response, increase on average within and across periods, particularly in the Hawkish populations.*

5 Discussion

Evolutionary game theory, like competitive equilibrium theory, predicts outcomes that are the “result of human action but not ... human design” (Ferguson, 1767). However, due to conceptual and technical difficulties, there have been few laboratory tests to date of the distinctive predictions of evolutionary game theory.

In this paper we introduce new laboratory techniques for playing population games in continuous time, and we obtain some striking results. In the one-population treatment, the evolutionary game prediction—that the fraction of Hawk play will approximate the symmetric mixed NE value of $2/3$ —turns out to be very accurate after the first few seconds of our two minute periods. In the two-population treatment, the overall rate of Hawk play is substantially lower, though not quite as low as predicted. More importantly, we indeed get divergent rates of Hawk play across the two populations, and the rates move towards the predicted extremes of the asymmetric pure NE. Moreover, we find that initial conditions (especially as proxied by final behavior in the previous period, but also by current random initialization) help predict which of the two pure NE is selected. The differential inclusion

specification helps identify when initializations will have greatest predictive power.

The experiment also sheds some light on issues neglected by evolutionary game theory. The underlying individual behavior shows an increasing degree of specialization among strategically identical players, especially in the two-population treatment and among Hawks. The dynamics seem to be driven, at least in part, by an interplay between the best response rates and incentives to best respond. In the one-population treatment, the payoffs quickly equalize between the two strategies and there is little incentive to best respond. By contrast, in the two-population treatment, the incentive to best respond increases as the state gets closer to an asymmetric pure NE, and the state gets closer to that NE as more players respond to the incentive.

An intriguing aspect of our experiment is that, even in later two-population periods, players don't immediately jump to where they left off in the previous period. Instead, they seem at first to engage in vendettas, as witnessed by the highest overall rates of Hawk play. Nevertheless, the asymmetry of Hawk play increases steadily throughout the period, and in the long run (of 120 seconds!) the "rationality" of pure NE prevails, despite its extreme inequity. Are there treatments that prolong the vendetta epoch? Smaller populations, perhaps, or devices that evoke an Us vs Them mentality? Can the two populations then work out a mutual accommodation, e.g., by alternating the two pure NE within or across periods? Will vendettas would disappear altogether in larger populations? These questions seem ripe for future research.

References

- Aubin, Jean-Pierre and Arrigo Cellina.** 1984. *Differential Inclusions*. Springer.
- Björnerstedt, Jonas and Jörgen Weibull.** 1996. "Nash equilibrium and evolution by imitation." In Arrow, K. J. et al., *The Rational Foundations of Economics Behavior*, St. Martin's Press, 155-81.
- Cooper, Russell and Douglas DeJong and Robert Forsythe and Thomas Ross.** 1989. "Communication in the battle of the sexes game: some experimental results." *The RAND Journal of Economics*, 20(4):568-587.
- Cooper, Russell and Douglas DeJong and Robert Forsythe and Thomas Ross.** 1993. "Forward Induction in the Battle-of-the-Sexes Games." *American Economic Review*, 83(5):1303-1316.
- Crawford, Vincent.** 1991. "An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination." *Games and Economic Behavior*, 3(1):25-29.
- Duffy, John and Nick Feltovich.** 2010. "Correlated Equilibria, Good and Bad: An Experimental Study." *International Economic Review*, Forthcoming.
- Friedman, Daniel.** 1991. "Evolutionary Games in Economics." *Econometrica*, 59: 637-666.
- Friedman, Daniel.** 1996. "Equilibrium in Evolutionary Games: Some Experimental Results." *Economic Journal*, 106: 1-25.
- Friedman, Daniel and Ryan Oprea.** 2010. "A Continuous Dilemma." *Mimeo*.
- Ferguson, Adam.** 1767 (2007). *An Essay on the History of Civil Society*. Cambridge University Press.
- Gale, John and Kenneth Binmore and Larry Samuelson.** 1995. "Learning to be Imperfect: The Ultimatum Game." *Games and Economic Behavior*, 8(1):56-90.
- Güth, Werner and M. Yaari** 1992. "An Evolutionary Approach to Explain Reciprocal Behavior in a Simple Strategic Game." in U. Witt. *Explaining Process and Change - Approaches to Evolutionary Economics*, 23-34.

- Harsanyi, John.** 1973. "Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points." *Int. J. Game Theory* 2:1-23.
- Hirsch, Morris and Stephen Smale.** 1974. *Differential Equations, Dynamical Systems, and Linear Algebra*. Academic Press.
- Huck, Steffen and Hans-Theo Normann and Jörg Oechssler.** 1999. "Learning in Cournot Oligopoly-An Experiment." *The Economic Journal*, 109(454):C80-C95.
- Huck, Steffen and Jörg Oechssler.** 1999. "The Indirect Evolutionary Approach to Explaining Fair Allocations." *Games and Economic Behavior*, 28(1):13-24.
- Maynard Smith, John.** 1982. *Evolution and the Theory of Games*. Cambridge University Press.
- Neugebauer, Tibor and Anders Poulsen and Arthur Schram.** 2008. "Fairness and reciprocity in the HawkDove Game." *Journal of Economic Behavior and Organization*, 66:234-250.
- Taylor, Peter and Leo Jonker.** 1978. "Evolutionary stable strategies and game dynamics." *Mathematical Biosciences*, 40(1-2):145-156.
- Schlag, Karl.** 1998. "Why Imitate, and if so, How? A Boundedly Rational Approach to Multi-Armed Bandits." *Journal of Economic Theory*, 18(1):130-156.
- Sandholm, William.** 2010. *Population games and Evolutionary Dynamics*. MIT Press.
- Smale, Stephen.** 1976. "Exchange processes with price adjustment." *Journal of Mathematical Economics*, 3: 211-226.
- Van Huyck, John and Raymond Battalio and Sondip Mathur and Patsy Van Huyck.** 1995. "On The Origin of Convention: Evidence From Symmetric Bargaining Games." *International Journal of Game Theory*, 24(2):187-212.
- Weibull, Jörgen.** 1995. *Evolutionary Game Theory*. MIT Press.

Appendix A: Mathematical Details

Sign preserving dynamics as differential inclusions

A *differential inclusion* on our state space is defined by a correspondence $V : [0, 1]^2 \rightrightarrows \mathfrak{R}^2$, $s = (s_1, s_2) \mapsto V(s) \subset \mathfrak{R}^2$. A *solution* to (or *trajectory* of) the differential inclusion V from a given initial condition $s_o = (s_{1o}, s_{2o})$ is a continuous map $T : [0, \infty) \rightarrow [0, 1]^2, t \mapsto s(t) = (s_1(t), s_2(t))$ such that $s(0) = s_o$ and, except perhaps for t in a measure zero subset of $[0, \infty)$, the tangent vector $\dot{s}(t)$ exists and satisfies $\dot{s}(t) \in V(s(t))$.

For sign preserving dynamics, the set $V(s)$ of allowable tangent vectors is (most of) the orthant defined by the signs of $(\Delta w_1, \Delta w_2)$. To write it out, use the sign function $\text{sgn}(x) = -$ if $x < 0$, $= +$ if $x > 0$, and $= o$ if $x = 0$, and let $\mathfrak{R}_{++} = \{(x, y) \in \mathfrak{R}^2 : x > 0, y > 0\}$ denote the open positive orthant, $\mathfrak{R}_{+-} = \{(x, y) \in \mathfrak{R}^2 : x > 0 > y\}$ denote the Southeast orthant, $\mathfrak{R}_{o-} = \{(x, y) \in \mathfrak{R}^2 : x = 0 > y\}$ denote the line heading South from the origin, etc. Also, let $\|\cdot\|$ denote the usual Euclidean norm, e.g., $\|\Delta w\|^2 = \Delta w_1^2 + \Delta w_2^2$. We impose a Lipschitz-Smale commensurability condition via the ε -annulus, $A^\varepsilon(s) = \{(x, y) \in \mathfrak{R}^2 : \varepsilon^{-1}\|\Delta w(s)\| \geq \|(x, y)\| \geq \varepsilon\|\Delta w(s)\|\}$.

Sign preserving dynamics can now be defined as solutions to a differential inclusion V such that, for some small $\varepsilon > 0$,

$$V(s) = A^\varepsilon(s) \cap \mathfrak{R}_{(\text{sgn}(\Delta w_1(s)), \text{sgn}(\Delta w_2(s)))} \quad (9)$$

for all $s \in (0, 1)^2$. On the boundaries $s_i = 0$, the set $V(s)$ of tangent vectors is the intersection of the closed half-plane $\dot{s}_i \geq 0$ with the RHS of (9); this prevents the share of Hawks from going negative. Likewise, $V(s)$ on the boundaries $s_i = 1$ intersects the RHS of (9) with the half-plane $\dot{s}_i \leq 0$ to prevent the share from exceeding 1.0.

It is routine (but a bit tedious) to verify that this correspondence V is what Sandholm (2010, Ch 6A) calls good UHC, and therefore has well-behaved solutions. Of course, the solutions are not unique—in some sense, that is the point of using differential inclusions rather than ODEs—but their tangent vectors can vary continuously (or even be constant) on the interior of regions where Δw_1 and Δw_2 have a given signs, e.g., on the transient and the absorbing rectangles in the Hawk-Dove game. The trajectory can have a kink, however, when it crosses

regions.

The Lipschitz-Smale annulus implies a positive lower bound on the velocity outside any neighborhood of the equilibrium points (where $\|\Delta w\| = 0$) and thus prevents the trajectory from stagnating in the interior of a rectangle or in the relative interior of an edge. This allows the intuitive geometric arguments in the text to be made rigorous in a straightforward manner. Consequently we have the following result.

Proposition 1 . *Let $\{s(t) : t \in [0, \infty)\}$ be a trajectory with initial condition $s(0) = s_o = (s_{1o}, s_{2o})$ for sign preserving dynamics of the HD game in Table 1. Then*

1. *There is some $\hat{t} > 0$ such that $s_i(t) \geq \frac{2}{3} \geq s_{-i}(t) \quad \forall t \in [\hat{t}, \infty)$ either for $i = 1$ or for $i = 2$.*
2. *If $s_{io} > \frac{2}{3} > s_{-io}$ then $s_i(t) > \frac{2}{3} > s_{-i}(t) \quad \forall t \in [0, \infty)$.*
3. *If $s_{io} > \frac{2}{3} > s_{-io}$ then $\lim_{t \rightarrow \infty} s_i(t) = 1$ and $\lim_{t \rightarrow \infty} s_{-i}(t) = 0$.*

Part 1 says that the open SW and NE rectangles are transient, and part 2 says that the open SE and NW rectangles are absorbing. Part 3 says that all trajectories starting in the SE (or NW) rectangle converge asymptotically to PNE1 (or to PNE2, respectively).

Three further remarks may be in order.

- Replicator dynamics and sign preserving dynamics offer contrasting predictions of what happens in the HD game near the unstable corners $(0, 0)$ and $(1, 1)$ of the state space. Replicator dynamics predict very slow adjustment, while sign preserving dynamics suggest adjustment at least as rapid there as anywhere else in the transient rectangles. We didn't test this prediction because our experiment provides almost no relevant data, and also because the contrast arises not from the basic nature of ODEs versus differential inclusions, but rather from their implementation in terms of growth rates (for replicator) versus rates of change (for the differential inclusions).
- Sign preserving dynamics do not always produce such clear predictions, even in 2x2 bimatrix games. For example, in a typical matching pennies type game, sign preserving dynamics predict only that trajectories will spiral clockwise (or counterclockwise,

depending on the parametrization). They allow the interior NE to be asymptotically stable (inward spirals), or neutrally stable (closed loops), or unstable (outward spirals). They even allow self-intersecting trajectories.

- With more than two alternative strategies available to some populations, sign preserving dynamics admit many different generalizations. As noted in Weibull (1995), Sandholm (2010) and elsewhere, these include payoff-monotone dynamics, payoff-positive dynamics, best-response dynamics, and many more.

Stability Analysis

Here we use standard linearization techniques (e.g., Hirsch and Smale, 1974) to determine the stability of the HD steady states under replicator dynamics.

Recall that for the given HD matrix, replicator dynamics are given by the following coupled pair of ordinary differential equations:

$$\dot{s}_1 = 3s_1(1 - s_1)(2 - 3s_2) = 6s_1 - 6s_1^2 - 9s_1s_2 + 9s_1^2s_2, \quad (10)$$

$$\dot{s}_2 = 3s_2(1 - s_2)(2 - 3s_1) = 6s_2 - 6s_2^2 - 9s_1s_2 + 9s_2^2s_1. \quad (11)$$

The Jacobian matrix for this system is:

$$J(s_1, s_2) = \begin{pmatrix} 6 - 12s_1 - 9s_2 + 18s_1s_2 & -9s_1 + 9s_1^2 \\ -9s_2 + 9s_2^2 & 6 - 12s_2 - 9s_1 + 18s_1s_2 \end{pmatrix} \quad (12)$$

Evaluating at the steady state PNE1, where $(s_1, s_2) = (1, 0)$, we have

$$J = \begin{pmatrix} -6 & 0 \\ 0 & -3 \end{pmatrix} \quad (13)$$

Here the eigenvalues are $\lambda_1, \lambda_2 = -6, -3 < 0$ and hence we have a sink, i.e., a locally asymptotically stable hyperbolic critical point.

The eigenvalues for the other asymmetric PNE2, where $(s_1, s_2) = (0, 1)$, are similarly seen to be $\lambda_1, \lambda_2 = -3, -6 < 0$, so PNE2 is another sink.

At the interior MNE steady state $(s_1, s_2) = (\frac{2}{3}, \frac{2}{3})$,

$$J = \begin{pmatrix} 0 & -2 \\ -2 & 0 \end{pmatrix}. \quad (14)$$

The characteristic equation now is $\lambda^2 - 4 = 0$ and the eigenvalues are $\lambda_1, \lambda_2 = 2, -2$. Since $\lambda_1 > 0 > \lambda_2$ we have a saddle point.

At the corner $(s_1, s_2) = (0, 0)$

$$J = \begin{pmatrix} 6 & 0 \\ 0 & 6 \end{pmatrix} \quad (15)$$

Here $\lambda_1 = \lambda_2 = 6 > 0$ which implies $(s_1, s_2) = (0, 0)$ is a source, i.e., a hyperbolic critical point for which all trajectories starting nearby exit any small neighborhood.

At the opposite corner $(s_1, s_2) = (1, 1)$

$$J = \begin{pmatrix} 3 & 0 \\ 0 & 3 \end{pmatrix} \quad (16)$$

Therefore $\lambda_1 = \lambda_2 = 3 > 0$ so here we have another source.

Finite Populations

Here we show that, with finite populations, the equilibrium states are a bit different than with infinite populations, but the modifications necessary for our experiment turn out to be quite minor.

To begin, suppose there is single finite population of n players. Then the fraction of Hawk players is $s_h = \frac{k}{n}$ for some $k \in \{0, 1, \dots, n\}$. The protocol is to match each subject with every other subject, but not with himself. Hence a Hawk player faces $k - 1$ Hawks among the $n - 1$ matches, and so the relevant state is $s = \frac{s_h n - 1}{n - 1}$. The indifference condition for a Hawk thus is $0 = \Delta w(s) = 6 - 9s$ or $\frac{2}{3} = s = \frac{s_h n - 1}{n - 1}$. Solving, we see that Hawks are indifferent when

$$s_h^* = \frac{2}{3} + \frac{1}{3n} \quad (17)$$

and that they have no incentive to switch when $s_h \leq s_h^*$.

Similarly, a Dove faces the fraction $\frac{s_d n - 1}{n - 1}$ of Doves, where $s_d = \frac{n - k}{n} = 1 - s_h$ is the overall fraction, including himself. Solving the indifference condition $\frac{s_d n - 1}{n - 1} = \frac{1}{3}$ yields

$$s_d^* = \frac{1}{3} + \frac{2}{3n} \quad (18)$$

as the state at which Doves are indifferent. Again, Doves are optimizing as long as $s_d \leq s_d^*$, or $s_h = 1 - s_d \geq 1 - s_d^*$.

Thus we have a Nash equilibrium of the finite single population game at any state $s_h \in [1 - s_d^*, s_h^*]$. Of course, in the limit as $n \rightarrow \infty$, the NE interval collapses to a single point $s^* = \frac{2}{3}$.

One session in our experiment had $n = 10$ subjects and the others had $n = 12$. In the the 10 player case, equations (17-18) yields the NE interval [.6,.7]. Thus any profile with 6 or 7 Hawks is a NE of the 10 player game. In the 12 player case, the NE interval is [.611,.694]. The only NE profiles are those with exactly 8 of the 12 players choosing Hawk, since 7/12 and 9/12 lie outside the NE interval.

Now consider the case of two finite populations of equal size n . Since each subject is matched with only subjects in the other population, the asymmetric pure NE are unchanged relative to the infinite population case. With $n = 6$ as in most of our two population sessions, the interior NE also remains unchanged, with 4 Hawks in each population. With $n = 5$, the interior NE $(s_1, s_2) = (\frac{2}{3}, \frac{2}{3})$ can not be supported by a pure strategy combination, since $\frac{3}{5} < \frac{2}{3} < \frac{4}{5}$, but it can be supported if at least one player mixes. For example, in NE each population could have 3 Hawk players plus 1 Dove player plus 1 player mixing 1/3 Hawk and 2/3 Dove.

Gini Coefficient

We now calculate the Gini coefficient when two equal-size populations play a PNE. Recall that the Gini coefficient for the MNE is zero since all subjects receive the same payout.

The Gini coefficient is defined as $\frac{A}{A+B}$ where A is defined as the area between the diagonal and the Lorenz curve and B is defined as the area below the Lorenz Curve. Since $A + B = \frac{1}{2}$,

we can write that Gini coefficient is simply $2A$.

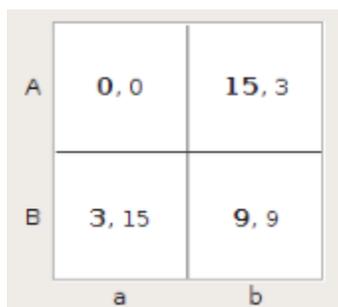
In the PNE, one population will receive a per capita payoff of 15 (Hawk) while the other population will receive 3 (Dove). Therefore the total is $3+15=18$, with one population receiving $\frac{3}{18} = \frac{1}{6}$ and the other $\frac{15}{18} = \frac{5}{6}$ of the total payoff. Hence the Lorenz curve consists of the two line segments connecting $(0, 0)$ to $(\frac{1}{2}, \frac{1}{6})$ and $(\frac{1}{2}, \frac{1}{6})$ to $(1, 1)$. It is now straightforward to calculate that $A = \frac{1}{6}$, and so the Gini coefficient at a PNE is $2A = \frac{1}{3}$.

Appendix B: Experimental Instructions

Instructions (C)

Welcome. This is an experiment in the economics of decision-making. If you pay close attention to these instructions you may earn a significant amount of money that will be paid to you in cash at the end of the experiment.

The Basic Idea



A	0, 0	15, 3
B	3, 15	9, 9
	a	b

In each of several periods, you will be able to choose one of two actions: A or B. Each period you will be matched with other players. Your earnings depend on the combination of your action and the other players actions that period.

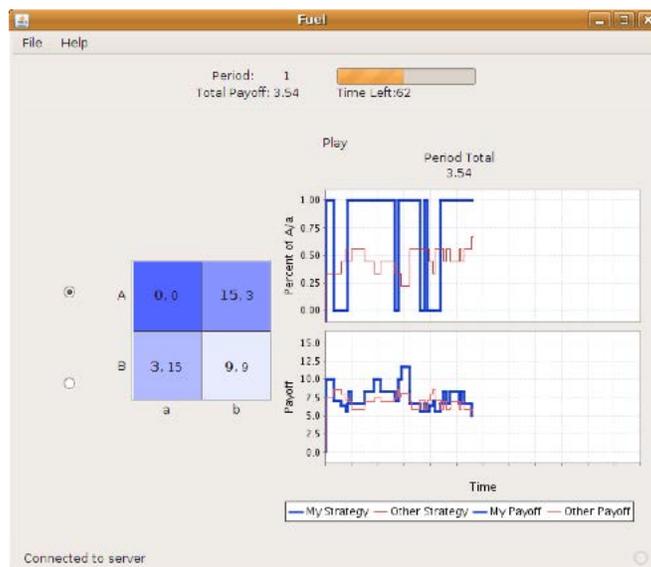
The earnings possibilities will be represented in a GAME MATRIX like the one above. Your action will determine the row of the matrix (A or B) and each other players action will determine a column of the matrix (a or b). The cell corresponding to this combination of actions will determine your EARNINGS. In each cell are two numbers. The first of the two numbers (shown in bold) is your earnings from this action combination. The second is the other player's earnings. You earn points from each match, and the points are scaled down by the number of other players.

For example, if there are 7 other players and 4 of them chose A and 3 chose B, then your payoff would be $(4*0 + 3*15)/7 = 45/7 = 6.43$ if you chose A, and it would be $(4*3 + 3*9)/7 = 39/7 = 5.57$ if you chose B.

You will not have to do this arithmetic yourself. The computer does the calculations and, as explained below, the bottom graph on your screen will display your earnings as you go along.

How to Play

There will be several periods. Each period will last 120 seconds and a counter at the top of the screen will show how much time is left. The computer randomly chooses the initial action, but



you can change your action at any time by clicking the two radio buttons or by using the up and down arrow keys. The row corresponding to your chosen action be highlighted in blue as in the figure, and the columns will be shaded in blue according to the number of players currently choosing that action. You and the other players may change your actions as often as you like each period.

The numbers in the payoff matrix are the payoffs you would earn if you maintained the same action throughout the period. For instance if you played B for the entire period and all other players played b in the example above, then you would earn 9 points and the other players also would earn 9 points each.

If you played A for the first half of the period and B for the second half while the other players played b for the entire period, your earnings would be $\frac{1}{2}(15) + \frac{1}{2}(9) = 12$, while the other players earnings would be $\frac{1}{2}(3) + \frac{1}{2}(9) = 6$. This is because you spent half of the period in the upper right corner and half in the lower right corner of the payoff matrix.

In general, your payoffs in the period will depend on how much time is spent in each of the

cells on the payoff matrix. The more time you spend in any one cell, the closer the final payoffs will be to the payoffs in that cell.

To the right of the screen are two graphs showing outcomes over the course of the period. The top graph shows your action (in blue) and the average action of all other players (in red) over the period. The graph is labeled Percentage of A. If this now reads 100 it means that at the moment you chose A. If it is 0 it means at that moment you chose B, and it switches between 0 and 100 as you switch actions.

The bottom graph shows your earnings over the course of the period in blue. The more area below your earnings curve, the more you have earned. In other words, the higher the blue line the more you are currently earning. The red line shows the corresponding average payoff for the other players.

Earnings

You will be paid at the end of the experiment based on the sum of point earnings throughout the experiment. These total earnings are displayed as the Total Payoff at the top of the screen.

Frequently Asked Questions

Q1. Is this some kind of psychology experiment with an agenda you haven't told us?

Answer. No. It is an economics experiment. If we do anything deceptive or don't pay you cash as described then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are meant to clarify how you earn money, and our interest is in seeing how people make decisions.

Q2. If I choose the rows and the other players chooses the columns, does their screen look different than mine?

Answer. On everyone's screen, the same choices are shown as rows. For example if another player chooses row B then it shows up on your screen as a choice of column b. Of course, the payoff numbers for any choice combination are the same on both screens, but are shown in a different place.

Q3. Who am I matched with? Everyone else in the room?

Answer. Sometimes you are matched with all other players in the room. Sometimes we divide the players into two groups and we match you only with players in the other group.