Strategic interactions and belief formation: An experiment

Kyle Hyndman* Antoine Terracol† Jonathan Vaksmann‡

March 2, 2009

Abstract

Traditional models of belief formation in repeated games assume adaptive players who do not take strategic interactions into account. We find that these approaches are limited in the sense that people think more strategically and realize that, in contrast with the classical view, their own actions are likely to influence their opponents' behaviour.

JEL codes: C91, D43, D83
Keywords: Game theory, Learning, Beliefs, Experiment

1 Introduction and motivations

In recent years, economists have studied ways to describe how people learn in a repeated game. To tackle the strategic complexity inherent in repeated games, belief learning models assume that players see their opponents' behaviour as generated by an exogenous process. Hence, typical proxies for beliefs are formed on the basis of the historical data of their opponents' play only and, according to this view, players do not realize that their own actions could influence their opponents' behaviour (*i.e.,* strategic interactions are neglected).

In this study, we run an experiment where subjects play repeated games under various conditions which are likely to have an impact on strategic considerations that drive

*Department of Economics, Southern Methodist University, 3300 Dyer Street, 301R, Dallas, TX 75275, hyndman@smu.edu

†EQUIPPE, Université de Lille, and Centre d'Économie de la Sorbonne, Université Paris 1 - Panthéon Sorbonne, CNRS, terracol@univ-paris1.fr

‡Centre d'Économie de la Sorbonne, Université Paris 1 - Panthéon Sorbonne, CNRS, jonathan.vaksmann@univ-paris1.fr

players' behaviour. We find that when players have relatively high incentives to behave strategically, their reasoning is more thoughtful than widely assumed.

2 Experimental design and procedures

In order to examine the role of strategic interactions in repeated games, we conducted four experiments. Inexperienced subjects were brought into the experimental laboratory at the University of Paris 1 Panthéon-Sorbonne[1] and were asked to play one of the games in Table 1 in fixed pairs for a total of 20 periods. Subjects were randomly assigned the role of either a row (r) or a column (c) player and were told that they would remain in that role for the duration of the experiment. Payoffs were denominated in experimental currency units and were converted into Euros at the conclusion of the experiment. Subjects earned, on average, €14.1 for their participation. In addition to written instructions (available upon request), subjects received an oral summary of the instructions.

All of our games have two pure-strategy Nash equilibria and one mixed-strategy Nash equilibrium. In these games, the pure strategy equilibria are Pareto rankable, both players strictly prefer the equilibrium (X,X) to the equilibrium (Y,Y). The mixed strategy equilibrium was $\{(0.8,0.2);(0.8,0.2)\}$.

Table 1: Payoff Matrices Used In The Experiments

HL

	X	Y
X	40, 45	8, 37
Y	39, 0	12, 32

HH

	X	Y
X	40, 45	0, 37
Y	37, 0	12, 32

LL

	X	Y
X	20, 45	8, 37
Y	19, 0	12, 32

LH

	X	Y
X	20, 45	0, 37
Y	17, 0	12, 32

One of the main reasons for neglecting strategic interactions in repeated games is that such behaviour is cognitively demanding. We conjecture that, when given sufficiently strong incentives, players are more likely to take strategic interactions into account. We designed our games to vary players' tendency to think more strategically. In order to measure this we vary parameters which are likely to play an important role. First, we vary the costs of not best-responding. Second, we vary the gains from convergence to the Pareto efficient equilibrium.

[1] The experiment was programmed using 'Regate' (Zeiliger, 2000).

Denoting $\pi_i(a,a')$, $i \in \{r,c\}$, player i's payoff when he plays a and his opponent plays a' and $E_i^a(p)$ player i's expected payoff from taking action a, $a \in \{X,Y\}$, given a belief of p, with p^* being the equilibrium mixing probability, the cost of not best-responding is:

$$E_i^X(p) - E_i^Y(p) = \theta_i \cdot (p - p^*),$$

with

$$\theta_i = \pi_i(X,X) - \pi_i(Y,X) + \pi_i(Y,Y) - \pi_i(X,Y).$$

The cost parameter, θ_i, is the optimization premium parameter of Battalio, Samuelson, and Van Huyck (2001). We conjecture that players will be more inclined to take strategic interactions into account carefully when the cost of not best-responding is relatively high.

Similarly, if the reward from coordinating on the Pareto efficient equilibrium is high, we would expect players to play strategically in order to "teach" their opponent (Camerer, Ho, and Chong, 2002, Ehrblatt, Hyndman, Ozbay, and Schotter, 2009, Terracol and Vaksmann, 2009) and thus to pay greater attention to the way he reacts. Therefore, in our experiments, we also vary the gain from moving from the inefficient to the efficient equilibrium. More precisely, for player i, $i \in \{r,c\}$, these gains can be written as:

$$\psi_i = \frac{\pi_i(X,X) - \pi_i(Y,Y)}{\pi_i(Y,Y)}.$$

We have four games according to the size (High or Low) of the expected cost and the gains of the beneficial action. In Table 1, the first letter above each game refers to the size of the benefits of the beneficial action while the second stands for the size of its cost. Notice that in all of our games the column player's payoffs are identical, while the row player's payoffs vary according to the costs and benefits of choosing action X. Moreover, the column player always has strictly weaker incentives than the row player.[2]

In this study we must elicit players' beliefs in order to perform a detailed examination of strategic thinking. In each round, before choosing their action, subjects reported their beliefs about the likely action of their match in that round. That is, they reported a probability vector, $b = (b_X, b_Y)$, where b_a represents the belief held by the subject associated to the action a of his opponent ($a \in \{X,Y\}$). A player's payoff when he reports b and his opponent actually uses action a is given by the following quadratic scoring rule:[3]

$$\left[8 - 4\left((1-b_a)^2 + b_{a'}^2\right)\right], \; a \neq a'.$$

At the end of each round, subjects were informed about the actions and stage game payoffs of both players.

[2] In games HL, HH, LL and LH we have respectively 34, 32, 38 and 30 subjects. ψ_r is 2.33 in games HL and HH and 0.67 in games LL and LH. θ_r is 5 in games HL and LL and 15 in games HH and LH. ψ_c and θ_c are respectively 0.41 and 40 in all games.

[3] See Nyarko and Schotter (2002) for a detailed description of this procedure.

3 Results

The aim of our study is to test whether players think strategically and anticipate their opponents' reactions to their own actions. Because beliefs may also depend on the history of the opponents' past actions, as postulated by traditional proxies used to describe players' belief-formation process, one must filter out the impact of these past actions to avoid spurious correlations between $a_i(t-1)$, player i's own action in the previous round, and his current beliefs. We thus examine the differences between stated beliefs and beliefs based on the past history of the game.

Adopt the terminology of Nyarko and Schotter (2002) and refer to beliefs based only on the history of the opponents' actions as "empirical" beliefs. Denote empirical beliefs by $\tilde{B}_i^a(t)$ and stated beliefs by $B_i^a(t)$. Next define $D_i^a(t) = B_i^a(t) - \tilde{B}_i^a(t)$ to be the difference between stated and empirical beliefs. Observe that since $\tilde{B}_i^a(t)$ is conditional on the history of the opponents' past actions, but not on the action chosen by player i in period $t-1$ (*i.e.,* $a_i(t-1)$), then if $D_i^a(t)$ depends on $a_i(t-1)$, so too must stated beliefs. In this case, we may conclude that players realize the influence of their own actions on their opponents' behaviour. In other words, players would take strategic interactions into account. We chose to model empirical beliefs with the γ-weighted beliefs model of Cheung and Friedman (1997), where the belief held by player i about the probability that player j will play action a in round $t+1$ is given by:

$$\tilde{B}_i^a(t+1) = \frac{\mathbb{1}_{\left(a_j(t)=a\right)} + \sum_{u=1}^{t-1} \gamma^u \mathbb{1}_{\left(a_j(t-u)=a\right)}}{1 + \sum_{u=1}^{t-1} \gamma^u} \tag{1}$$

where $\mathbb{1}_{\left(a_j(t)=a\right)}$ equals one if player j has played action a in round t, and zero otherwise. Actions played in a given round are discounted with time at rate $\gamma \in [0,1]$. According to these empirical proxies for beliefs, players form conjectures only on the basis of the history of their opponent's actions; hence, strategic interactions do not play any role.

We estimate the model of equation (1) at the individual level, and compute $\hat{D}_i^a(t)$, as the difference between stated beliefs and estimated empirical beliefs in round t.

We then examine whether $\hat{D}_i^a(t)$ can be explained by the action taken by player i in the previous round — $a_i(t-1)$.[4] If so, this would indicate that players (at least partly) base their actions on motivations beyond those suggested by classical adaptive proxies. For this reason, we refer to it as a *sophistication bias*.

We first display a graphical analysis that can be complemented with paired t-tests. Figure 1 shows, for each game and each type, the mean value of $\hat{D}_i^X(t)$ according to whether or not X was played in the previous round.

It is clear from Figure 1 that the $\hat{D}_i^X(t)$ are almost always higher when X rather than Y has been played in the previous round. Moreover, the residuals tend to be positive when

[4]Note that $\hat{D}_i^Y(t) = -\hat{D}_i^X(t)$, so our conclusions are valid across actions.

Figure 1: $\hat{D}_i^X$

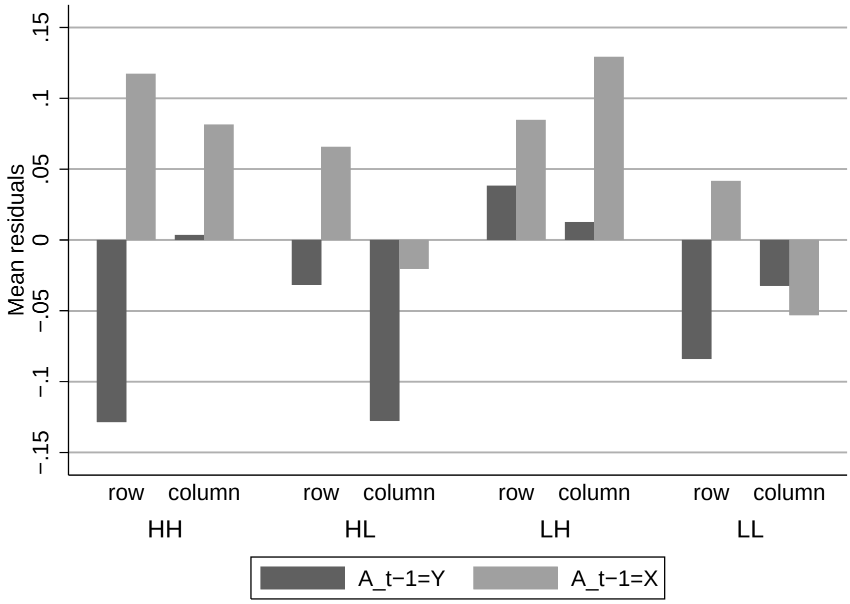

$a_i(t-1) = X$ and negative when $a_i(t-1) = Y$. This is consistent with the hypothesis that players take their own actions into account when forming beliefs.

We also run a series of paired t-tests on the difference of the mean $\hat{D}_i^X(t)$ at the individual level. In other words, we test whether, when one controls for the past history of the opponent's actions, beliefs towards X vary significantly according to the player's previous action. The results are displayed in Table 2. The first line shows the mean differences between $\hat{D}_i^X(t|a_i(t-1) = X)$ and $\hat{D}_i^X(t|a_i(t-1) = Y)$. The following two lines give the p-value and the number of observations, respectively.[5]

Table 2: Paired t-tests on $\hat{D}_i^X(t)$, according to previous action

	HL		*HH*		*LL*		*LH*	
	row	column	row	column	row	column	row	column
difference	0.097	0.107	0.245	0.077	0.125	-0.021	0.046	0.116
p-value	0.087	0.038	0.013	0.187	0.061	0.736	0.636	0.255
N	16	15	12	12	17	16	14	14

These results show that row players' beliefs vary significantly according to his previous action in all games except game *LH* where his incentives to pay attention to strategic

[5] The number of observations varies within games since some players never played X or never played Y. These players are thus excluded from this analysis.

interactions are relatively weak. For column players, who always have particularly low incentives to think strategically, the differences are broadly insignificant. Only in the game *HL* is this difference significant. This is somehow surprising since column players' incentives remain unchanged across games. This suggests a common understanding of players' incentives to think strategically.

Does the sophistication bias vary according to the gains and cost of the beneficial action? If so, are one's own gains and costs the only parameters that matter, or do the gains and costs of the opponent also influence the way players form their beliefs?

To answer this question, we estimate fixed-effect regressions of the residuals on the previous action, and the previous action interacted with the gain and cost parameters of both players:

$$D_i^X(t) = \beta_0 + \left[\beta_1 + \beta_2\psi_i + \beta_3\theta_i + \beta_4\psi_j + \beta_5\theta_j\right] \times \mathbb{1}_{(a_i(t-1)=X)} + \nu_i + \varepsilon_{i,t}$$

for $t > 1$; $i, j \in \{r, c\}$; $j \neq i$.

The results are collected in Table 3. We estimate four types of models. In model (1), coefficients β_2 to β_5 are constrained to zero. In model (2), coefficients β_4 and β_5 are constrained to zero (*i.e.,* only one's own ψ and θ matter). In model (3), coefficients β_2 and β_3 are constrained to zero (*i.e.,* only one's opponent's ψ and θ matter). Finally, model (4) is the unrestricted model. We use the gain and cost parameters in deviation from their sample mean, and estimate all four models on pooled data, then models (1) and (2) on row players only, and models (1) and (3) on column players only.[6]

In Table 3, β_1 gives the average sophistication bias. Coefficients β_2 to β_5 indicate how this bias varies with the gains and costs. We hypothesize that players pay more attention to the way their opponent forms their beliefs when the stakes are high. If this is true, then stated beliefs should be further away from empirical beliefs as one's ψ and θ increase. We thus expect $\beta_2 > 0$ and $\beta_3 > 0$. Similarly, players should think their opponents react more to their actions when the opponent's ψ and θ are high. Thus, we expect β_4 and β_5 to be positive as well.

The results indicate that players are responsive to incentives when forming beliefs. While the average impact of having played X in the previous period is always positive and significant, it is modified by the size of the gains and costs parameters. Interestingly, the impact of players' benefits of efficient coordination on the sophistication bias is more salient than the impact of its cost. Indeed, with the exception of model (3) on pooled data, $\hat{\beta}_3$ and $\hat{\beta}_5$ remain insignificant in all relevant cases while $\hat{\beta}_2$ and $\hat{\beta}_4$ are always significantly positive.

[6]The full set of models are not estimated on rows and columns only because ψ_i and θ_i (resp. ψ_j and θ_j) vary only across row players (resp. column players).

4 Conclusion

Usual proxies used in learning models to describe players' belief-formation process postulate that people do not take strategic interactions into account, thus neglecting a concept at the heart of game theory. To investigate its limits, we use experimental games to elicit players' true beliefs and compare them to a common general model of belief-formation. We find that belief proxies are biased as players take strategic interactions into account, particularly when incentives to play strategically are high. Hence, we conclude that the common assumption of purely adaptive players is not innocuous.

References

BATTALIO, R., L. SAMUELSON, AND J. VAN HUYCK (2001): "Optimization Incentives and Coordination Failure in Laboratory Stag Hunt Games," *Econometrica*, 69(3), 749–64.

CAMERER, C., T.-H. HO, AND J.-K. CHONG (2002): "Sophisticated ewa learning and strategic teaching in repeated games," *Journal of Economic Theory*, (104), 137–188.

CHEUNG, Y.-W., AND D. FRIEDMAN (1997): "Individual learning in normal form games: Some laboratory results," *Games and Economic Behavior*, 19(1), 46–76.

EHRBLATT, W. Z., K. HYNDMAN, E. Y. OZBAY, A. SCHOTTER (2009): "Convergence: An experimental study of teaching and learning in repeated games," mimeo.

NYARKO, Y., AND A. SCHOTTER (2002): "An Experimental Study of Belief Learning using Elicited Beliefs," *Econometrica*, 70(3), 971–1006.

TERRACOL, A., AND J. VAKSMANN (2009): "Dumbing down rational players: Learning and teaching in an experimental game," *Journal of Economic Behavior and Organization* , in press.

ZEILIGER, R. (2000): "A presentation of regate, internet based software for experimental economics.," GATE, `http://www.gate.cnrs.fr/~zeiliger/regate/regateintro.ppt`.

Table 3: Results

	Pooled				Row Players		Column Players	
	(1)	(2)	(3)	(4)	(1)	(2)	(1)	(3)
β_0	-0.052^{***} (0.009)	-0.052^{***} (0.009)	-0.052^{***} (0.009)	-0.052^{***} (0.009)	-0.061^{***} (0.013)	-0.061^{***} (0.013)	-0.043^{***} (0.013)	-0.042^{***} (0.013)
β_1	0.083^{***} (0.013)	0.081^{***} (0.013)	0.081^{***} (0.013)	0.081^{***} (0.013)	0.104^{***} (0.018)	0.103^{***} (0.018)	0.060^{***} (0.020)	0.059^{***} (0.020)
β_2		0.061^{***} (0.022)		0.058^{**} (0.022)		0.057^{***} (0.022)		
β_3		0.001 (0.001)		0.003 (0.003)		0.004 (0.004)		
β_4			0.039^{*} (0.023)	0.039^{*} (0.023)				0.039 (0.024)
β_5			0.003^{**} (0.001)	0.004 (0.003)				0.003 (0.004)

Significance levels: $*$: 10%; $**$: 5%; $***$: 1%. Standard errors in parenthesis