



ELSEVIER

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb



Feedback spillover and analogy-based expectations: A multi-game experiment [☆]

Steffen Huck ^{a,b}, Philippe Jehiel ^{c,a,b,*}, Tom Rutter ^{a,b}

^a University College London, United Kingdom

^b ELSE, United Kingdom

^c Paris School of Economics, France

ARTICLE INFO

Article history:

Received 13 July 2009

Available online xxxx

JEL classification:

C72

D82

Keywords:

Analogy-based expectation

Information processing

Experiments

Accessibility

Interactive learning

Feedback spillover

ABSTRACT

We consider a multi-game interactive learning environment in which subjects sometimes only have access to the aggregate distribution of play of the opponents over the various games and sometimes are told the joint distribution of actions and games in a more or less accessible way. Our main findings are: 1) In the presence of feedback spillover, long run behaviors stabilize to an analogy-based expectation equilibrium (Jehiel, 2005). 2) Faced with the same objective feedback, the long run behaviors are sometimes better described by Nash equilibrium and sometimes they are better described by the analogy-based expectation equilibrium depending on the accessibility of the feedback.

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

A number of solution concepts have been proposed to model the interaction of boundedly rational players in games, but relatively few of these concepts have been tested in the lab. In this paper, we test experimentally a solution concept introduced in Jehiel (2005) and called the analogy-based expectation equilibrium in which the bounded rationality of players is reflected by the fact that players have only coarse knowledge about the strategy of their opponent. Specifically, in contexts in which players may be called to play in different games (or nodes in extensive form games), Jehiel (2005) considers players who understand only the aggregate distribution of play of their opponent over various games, and who play a best-response as if the opponents behaved in each game according to the aggregate distribution.

The analogy-based expectation equilibrium – which is parameterized by how each player bundles the various games into analogy classes – is viewed as a steady state outcome of a (learning) process in which players would receive feedback about the aggregate distributions of their opponent's actions in their various analogy classes and nothing else (see Jehiel, 2005; Jehiel and Koessler, 2008; Ettinger and Jehiel, 2010).

[☆] We wish to thank seminar participants at NYU, Princeton, the “Computational Social System and the Internet” workshop at Dagstuhl 2007 as well as many colleagues for valuable comments. We also more especially thank Tom Palfrey, an associate editor and two referees for helpful comments that helped focus the paper better. We also wish to thank the Centre for Economic Learning and Social Evolution (ELSE) at UCL. Jehiel thanks the European Research Council for support.

* Corresponding author at: Paris School of Economics, 48 Bd Jourdan, 75014 Paris, France.

E-mail address: jehiel@enpc.fr (P. Jehiel).

Our experiment considers the simplest setup in which this theory has some bite, that is, a setup in which two two-person games A and B with the same action spaces for Column and Row players are being played. At the start of the experiment, populations of subjects are assigned to the role of Column or Row players. Each subject plays each of the games A and B thirty times in a random order. In each round, subjects from the population of Column players and subjects from the population of Row players are randomly matched to play either game A or B , and in each round there is an equal proportion of each game.

A key aspect of the experiment is the feedback provided to subjects between the various rounds. We consider four different treatments. In all treatments, the feedback concerns the distribution of play over the last five rounds of the population of subjects assigned to the opponent's role. In all treatments, Column players receive precise feedback about the game they are to play in the current round. That is, Column players who are to play game A are informed of the number of times the various actions of Row players were chosen when game A was played in the last five rounds. And similarly when they are to play game B .

The various treatments vary only with respect to the feedback provided to Row players. In treatment FINE, this feedback is similar to that provided to Column players. That is, when a Row subject is called to play game $X = A$ or B , this subject is being told the number of times the various actions of Column players were chosen in the last five rounds when game X was played. In treatment COARSE, the feedback is aggregated over the two games, that is, there is feedback spillover. More precisely, when a Row subject is called to play game $X = A$ or B , this subject is being told the total number of times the various actions of Column players were chosen in the last five rounds adding up the instances in which games A and B were played (there is no information as to whether the actions were made in game A or B). In treatments FEASIBLE and HARD, Row players receive feedback about the joint distribution of actions and the game in which the actions were chosen over the last five rounds. In both treatments agents with unlimited cognitive powers would be able to construct game-specific data, that is, they would have feedback as in FINE. With cognitive limitations, however, it may be more or less difficult to disentangle the provided information. It will turn out that while deriving the distribution of actions by games is possible in FEASIBLE, it is much harder in HARD (these points are elaborated in much more detail in the next section).

In addition to these four treatments, we have two further control treatments called SOLO in which subjects play only one game (60 times game A or 60 times game B) and obtain feedback about the distribution of actions of the subjects assigned to their opponent's role over the last five rounds.

In all treatments, subjects when called to play game X can see the payoff they would obtain for all profiles of actions. But, subjects are not told the corresponding payoffs for their opponent. Moreover, subjects are not informed of their performance until the very end of the experiment at which time they receive a payment that aggregates the payoffs over all sixty periods.

The games A and B chosen for the experiment are such that they admit a single Nash equilibrium and also a unique analogy-based expectation equilibrium in which the Row player bundles the two games into one analogy class and the Column player uses the finest analogy partition (this analogy-based expectation equilibrium is referred to as ABEE hereafter). The Nash equilibrium and the ABEE which both employ pure strategies correspond to distinct actions both in games A and B and both for the Column and the Row players, thereby making it easier to identify whether one or the other is being played. Finally, games A and B are dominance solvable, which implies that fictitious play dynamics applied to each game in isolation converge to the unique Nash equilibrium (see Nachbar, 1990; Milgrom and Roberts, 1991). Similarly, we show that an analogy-based fictitious play dynamics in which row subjects best-respond to the empirical distribution of column players's actions over the last five rounds and over the two games and column players play a best-response to the empirical distribution of row players' actions over the last five rounds in the game to be played in the current round converges to the ABEE.

Our findings are as follows. First, we observe that in all treatments the pattern of play stabilizes (roughly after 15 rounds). Second, in different treatments it stabilizes at different action profiles. In treatments SOLO, FINE, and FEASIBLE, behavior (mostly) converges to Nash equilibrium. In treatments COARSE and HARD behavior (mostly) converges to ABEE.

The fact that in SOLO we get convergence to Nash equilibrium is expected given that the games are dominance solvable.¹ The fact that in FINE we get similar behaviors as in SOLO is also expected given that there is little room for feedback spillover in FINE (somehow the feedback provided to subjects is exactly what subjects need to know to play the game, and the personal memory of the subject that also contains the feedback for the other game is of little relevance).

The finding for COARSE gives support to the implication of ABEE theory. A priori, there are several conjectures about column players' play that would be consistent with the coarse feedback received by row players (roughly any pair of distribution over games A and B whose average coincides with the aggregate distribution). Our experimental results suggest that, as required in ABEE, row players in COARSE are well described as playing a best-response to the conjecture that column players play according to the aggregate distribution in each of the games A and B .²

The findings for FEASIBLE and HARD reveal that by playing on the framing of the same objective feedback (here the joint distribution of actions and games) one can induce very different long run outcomes (Nash behavior and ABEE behavior,

¹ See for example Lieberman (1960).

² Put differently, these findings provide support to the ABEE as a relevant selection of self-confirming equilibrium when a player receives as feedback the aggregate distribution of the opponent's play over different games (see Battigalli, 1987; Dekel et al., 2004, for a definition of self-confirming equilibrium for arbitrary/unstructured feedback).

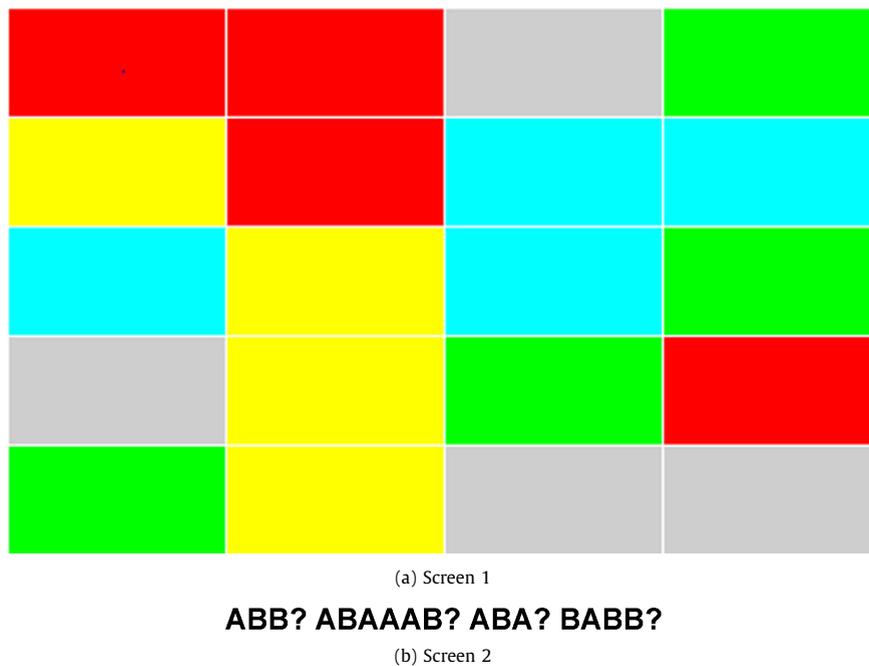


Fig. 1. HARD treatment.

respectively).³ The similarity between COARSE and HARD suggests that in HARD, Row subjects process information as if they were only informed of the aggregate distributions of behaviors over the two games. The similarity between FEASIBLE and FINE (and SOLO) reveal that in FEASIBLE row subjects are able to avoid the feedback spillover, and in the long run Nash equilibrium is played.

The rest of the paper is organized as follows. In the next section we present in more detail some aspects of the feedback used in FEASIBLE and HARD. In Section 3 we then move on to describe the analogy-based expectation equilibrium and apply it to the environment considered in our experiment. We also present convergence results for two (fictitious-play) learning models that help understand our experimental results better. In Section 4 we present further specifics of our experimental design and procedures. Section 5 contains the results of the experiments. Section 6 reviews the related literature and Section 7 concludes.

2. A closer look at the feedback structure

In both treatments, FEASIBLE and HARD, each column in the matrix game is given a color, and Row players are shown the distribution over the last five rounds of Column players' actions using colored boxes containing letters. The color of each box refers to a choice of column and the letters refer to the game (*A* or *B*) in which the corresponding column was chosen. The feedback in HARD is presented in two consecutive screens (with no permission to take notes or to go back to the previous screen). In the first screen, shown in Fig. 1(a), Row subjects can see the distribution of actions (represented as color boxes) over the last five rounds, i.e., they are shown twenty colored boxes (there were overall eight subjects by session, four being assigned to the role of Column players), each representing one past choice of a column. However, on that screen alone they cannot infer the type of game in which a particular action was chosen. This information is presented on the next screen, shown in Fig. 1(b), where subjects can see the corresponding distribution of games in a string of letter in which the first letter indicates the game played in the upper left corner of the first screen, the second letter indicates the game just below and so on.⁴ If subjects perfectly memorize the color pattern from screen 1, they can disentangle the distribution for both games from screen 2. If they have difficulties memorizing the precise pattern, they may, instead, just keep track of the aggregate color composition of screen 1 which represents the aggregate distribution over Column's choices in both games. In any case, this aggregate distribution is more easily accessible than the separate distributions are.⁵

³ Usually, framing effects are associated with one-shot decision problems and they are perceived to be temporary phenomena (see Tversky and Kahneman, 1981). In a different context, Huck et al. (2004) provide evidence that framing the interaction in economic or abstract terms may lead to more or less collusion in a repeated Cournot interaction.

⁴ The grey boxes in Fig. 1(a) and the question marks in Fig. 1(b) stand for those cases in the first five rounds in which some matches are missing.

⁵ See Higgins (1996) or Khaneman (2003) for an exposition of the accessibility idea.



Fig. 2. FEASIBLE treatment.

After the feedback screens are shown, Row subjects click to move to a subsequent screen where they are informed whether they are in game *A* or *B* (for this round) and they then have to choose an action. Again, once the player clicks to move forward, it is not possible to go back to the feedback screen so subjects have to remember what they saw on the feedback screens when choosing an action.

As our results summarized below will make clear, it appears that the patterns of behavior in *HARD* are similar to those in *COARSE* in which Row subjects get only to know the aggregate distribution of Column players' actions over the two games *A* and *B*. This is consistent with the interpretation that Row subjects in *HARD* mostly memorize the first feedback screen while paying little attention to the second feedback screen.

By contrast, in *FEASIBLE*, the feedback is presented in just one screen with the game *A* or *B* appearing in the center of the colored box (that stands for the action played in the corresponding game, see Fig. 2). As our results will make clear, it appears that the patterns of behaviors in *FEASIBLE* and *FINE* are more or less similar. This is consistent with the interpretation that Row subjects manage in *FEASIBLE* to play as if they have access to the fine distribution of actions by game.

3. Background and theory

We first describe the analogy-based expectation approach (in a setup appropriate for our experiment) and then apply it to the specific environment considered in our experiment. We next describe the corresponding learning models and provide some convergence results for these. In Appendix A, we develop the analogous models when players are assumed to use the quantal response model (as defined in McKelvey and Palfrey, 1995) rather than ordinary best-responses (the quantal response versions are used to analyze some of the data).

3.1. General background

Consider a family of normal form games where each game is denoted by $\omega \in \Omega$. Each game has two players i and j . For each ω , the action space of player i is A_i and the action space of player j is A_j . Action spaces A_i and A_j are finite. The payoff obtained by player i in game ω when $(a_i, a_j) \in A_i \times A_j$ is played is denoted by $u_i(a_i, a_j; \omega)$. The probability of game ω is denoted by $p(\omega)$. We assume that each player i knows which game $\omega \in \Omega$ he is playing (or at least i 's payoff in ω).

A strategy of player i is a mapping $\sigma_i : \Omega \rightarrow \Delta A_i$ where $\sigma_i(a_i | \omega)$ denotes the probability with which action $a_i \in A_i$ is chosen by player i in game ω .

Each player i is endowed with an analogy partition An_i over Ω . The element of An_i containing ω is denoted by $\alpha_i(\omega)$ and called the analogy class of player i at ω . Player i is assumed to understand only the aggregate behavior of player j in every analogy class in An_i . Formally, given the strategy σ_j of player j , the strategy of player j perceived by player i (given An_i) is defined by the function $\bar{\sigma}_j : \Omega \rightarrow \Delta A_j$ such that for all $\omega \in \Omega$ and $a_j \in A_j$

$$\bar{\sigma}_j(a_j | \omega) = \frac{\sum_{\omega' \in \alpha_i(\omega)} p(\omega') \sigma_j(a_j | \omega')}{\sum_{\omega' \in \alpha_i(\omega)} p(\omega')} = \sum_{\omega' \in \Omega} p(\omega' | \alpha_i(\omega)) \sigma_j(a_j | \omega') \tag{1}$$

Table 1
Normal form games used in the experiment.

A	a	b	c	d	e
α	25,10	0,10	10,20	0,0	0,0
β	20,15	15,0	5,0	10,0	0,0
γ	15,0	10,0	0,0	5,25	25,0

B	a	b	c	d	e
α	15,0	20,10	10,0	5,5	0,0
β	0,10	25,10	0,0	10,20	0,0
γ	10,0	15,0	5,25	0,0	25,0

That is, given the strategy σ_j of player j , player i perceives only the aggregate behavior of player j in each analogy class where the weight assigned to a specific game ω' of an analogy class is proportional to $p(\omega')$. In an analogy-based expectation equilibrium, for each $\omega \in \Omega$, each player i plays a best-response to how i perceives j is playing. That is,

Definition. A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is an analogy-based expectation equilibrium given the analogy partitions An_1, An_2 if for all $i, \omega \in \Omega$ and a_i^* in the support of $\sigma_i(\omega)$:

$$a_i^* \in \arg \max_{a_i \in A_i} \sum_{a_j \in A_j} \bar{\sigma}_j(a_j | \omega) u_i(a_i, a_j; \omega)$$

where $\bar{\sigma}_j(a_j | \omega)$ is given by (1).

Several comments are in order about this solution concept: 1) When all players use the finest analogy partition, the analogy-based expectation equilibrium coincides with Nash equilibrium. That is, in each game ω a Nash equilibrium is being played. 2) The analogy-based expectation equilibrium requires that player i knows $u_i(a_i, a_j; \omega)$ for all a_i, a_j and ω but not $u_j(a_i, a_j; \omega)$ as long as player i knows the aggregate distribution of play $\sum_{\omega \in \alpha_i} p(\omega) \sigma_j(a_j | \omega) / \sum_{\omega \in \alpha_i} p(\omega)$ in each analogy class α_i . This knowledge is viewed as the result of learning as explained below. 3) An analogy-based expectation equilibrium can be viewed as a selection of self-confirming equilibrium. It is a selection based on the idea that players consider the simplest (as opposed to any) conjecture consistent with their knowledge. Our experiment will help provide support to this selection.

3.2. Application to our experiment

In our experiment, we consider two games $\omega = A$ and B whose payoff matrices for the Row ($i = 1$) and Column player ($i = 2$) are depicted in Table 1. The two games A and B are played with the same frequency, i.e., $p(A) = p(B) = \frac{1}{2}$. Two profiles of analogy partitions will be relevant for our experiment. In each of them, the Column player, i.e. player 2, uses the fine analogy partition $An_2^f = \{\{A\}, \{B\}\}$. For the Row player, i.e. player 1, either he uses the fine partition, $An_1^f = \{\{A\}, \{B\}\}$, or he uses the coarse partition that mixes the two games A and B into one analogy class, $An_1^c = \{\{A, B\}\}$.

When the Row player's analogy partition is fine

When both players use the fine analogy partition, ABEE coincides with the Nash equilibrium of the complete information games A and B , as already noted. Games A and B are dominance solvable. Thus, they admit a single Nash equilibrium. This observation together with the equilibrium strategies are gathered in the following proposition.

Proposition 1. Games A and B are dominance solvable. The unique Nash equilibrium in game A is that Row player plays α , $\sigma_1(A) = \alpha$, and the Column player plays c , $\sigma_2(A) = c$. The unique Nash equilibrium in game B is that Row player plays β , $\sigma_1(B) = \beta$, and the Column player plays d , $\sigma_2(B) = d$. This is also the unique analogy-based expectation equilibrium when both players use the fine analogy partition.

Proof. In both games A and B action e is strictly dominated (by a mixture over a and d in A and a mixture of b and c in B , respectively) for the Column player. After eliminating action e action γ is strictly dominated for the Row player (by action β in A and action α in B). Following these eliminations, in game A , actions d and b are strictly dominated by a and a mixture of a and c , respectively. In game B , actions c and a are strictly dominated by b and a mixture of b and d , respectively. Finally, with the remaining actions, α in A and β in B strictly dominate β and α respectively, and we can conclude. \square

When the Row player's analogy partition is coarse

The following proposition characterizes the unique analogy-based expectation equilibrium when the column player uses the fine partition and the row player uses the coarse partition (this is called ABEE hereafter).

Proposition 2. Let $An_1 = An_1^c$ and $An_2 = An_2^f$ be the analogy partitions of players 1 and 2, respectively. There is a unique analogy-based expectation equilibrium (σ_1, σ_2) in which the Row player (player 1) plays β in game A and α in game B ($\sigma_1(A) = \beta, \sigma_1(B) = \alpha$) and the Column player (player 2) plays a in game A and b in game B ($\sigma_2(A) = a, \sigma_2(B) = b$).

Before we prove Proposition 2, observe that the strategies in ABEE are markedly different from the Nash equilibrium strategies. There is a complete swap between games A and B for the strategy of the row player, and the support of actions of the column player is $\{c, d\}$ in Nash equilibrium and it is $\{a, b\}$ in ABEE.⁶

Proof. It is easy to understand why the above strategy profile defines an analogy-based expectation equilibrium. The Column player plays a in A because this is the best-response to β ; she plays b in B because this is the best-response to α . The aggregate behavior of the Column player is a balanced mix of a and b (remember that $p(A) = p(B) = \frac{1}{2}$). Thus, $\bar{\sigma}_2(a | \omega) = \bar{\sigma}_2(b | \omega) = \frac{1}{2}$ for $\omega = A$ and B . Given the expectation that the Column player plays a and b with an equal frequency, the Row player finds it optimal to play β in game A (because $\frac{20+15}{2} > \frac{\max(20+0, 15+10)}{2}$) and α in game B (because $\frac{15+20}{2} > \frac{\max(0+25, 10+15)}{2}$).

It is easily checked that there is no other analogy-based expectation equilibrium. Indeed, action e is strictly dominated for the Column player in both games. After eliminating e , action γ for Row player is strictly dominated in both games. After eliminating γ , we can now eliminate actions d and b in A (they are dominated by a and any strictly convex combination of a and c , respectively) and c and a in B (they are dominated by b and any strictly convex combination of b and d , respectively). Finally, we use the analogy partition of row player to note that if the column player chooses any mixture of a and c in A and of b and d in B such that the row player perceives an equal mixture of both strategies in both games, then action β in A and action α in B dominates actions α and β , respectively (to see this, note that for all $\lambda \in (0, 1)$, $\frac{1}{2}(15\lambda + 10(1 - \lambda)) - \frac{5}{2} > 0$). The ABEE follows from the column player's best response to this. \square

Learning models

Jehiel (2005) and Jehiel and Koessler (2008) view the correctness of the perceptions $\bar{\sigma}$ in an analogy-based expectation equilibrium as the outcome of converging learning dynamics. We now define and analyze such learning processes that are guided by the design of our experiment.

In our experiment, each session of each treatment consists of several rounds $t = 1, 2, \dots$. In each round, four Column players and four Row players are randomly matched to play one of the games $\omega = A$ or B , and each game is played exactly twice, e.g. in two of the four matches $m = 1, 2, 3, 4$. We call M_ω^t the set of matches in which game $\omega = A$ or B is played in round t , and we refer to $\omega^t(m)$ as the game being played in match m of round t .

The feedback given to subjects concerns the behaviors of opponents (i.e., subjects assigned to the role of the other player) in the last five rounds.⁷ Two learning models are considered according to whether the feedback is about the aggregate behaviors in the two games or the behaviors game by game. For each game $\omega = A$ or B , we refer to

$$BR_i^\omega(x_j) = \arg \max_{a_i \in A_i} \sum_{a_j \in A_j} x_j(a_j) u_i(a_i, a_j; \omega)$$

as player i 's best response in game ω to the distribution $x_j \in \Delta A_j$ of player j 's actions where $x_j(a_j)$ denotes the weight assigned to action a_j in x_j .

For each player i , we let $a_i^{t,m}$ denote the action played by the player assigned to the role of player i in match m of round t . Given the belief $x_j(t+1)$ about player j 's behavior considered by player i in round $t+1$, our learning dynamics requires that for $t > 5$, and each $m = 1, 2, 3, 4$,

$$a_i^{t+1,m} \in BR_i^\omega(x_j(t+1))$$

where

$$\omega = \omega^{t+1}(m)$$

The two learning models differ only in how $x_j(t+1)$ is determined. In the fine feedback case, we have:

$$x_j^f(t+1) = \frac{1}{10} \sum_{k=t-4}^t \sum_{m' \in M_\omega^k} a_j^{k,m'} \tag{2}$$

That is, $x_j^f(t+1)$ is the empirical distribution of actions of subjects assigned to the role of player j over those matches in which game $\omega = \omega^{t+1}(m)$ (the one to be played) was played in the last five rounds.

⁶ Observe that the behavior of Column player is different in ABEE and Nash equilibrium even though this player uses the same analogy partition in the two cases. This is because Column player is best-responding to different distributions of Row player's action in ABEE and Nash equilibrium.

⁷ In the first five rounds, the feedback covers the entire history of play.

In the coarse feedback case, we have:

$$x_j^C(t+1) = \frac{1}{20} \sum_{k=t-4}^t \sum_{m'=1,2,3,4} a_j^{k,m'} \quad (3)$$

That is, $x_j^C(t+1)$ is the empirical distribution of actions of subjects assigned to the role of player j in all matches over the last five rounds.

The two learning models that we consider have the Column players following the fine feedback belief dynamics. In one learning model – that we call the Fine learning model – Row players follow the fine feedback belief dynamics. In the other – that we call the Coarse learning model – they follow the coarse feedback belief dynamics.

Observe that the Fine learning model is a variant of fictitious play dynamics with a five period window in which the learning in the two games A and B can be treated separately (there is no feedback spillover). The Coarse learning model can be viewed as the analog to fictitious play dynamics taking into account the feedback spillover between the two games. The following proposition states the convergence properties of these two learning models.

Proposition 3. *Whatever the initial conditions, play converges to the Nash equilibrium in the Fine learning model and to ABEE in the Coarse learning model.*

Proof. 1) *Convergence in Fine learning model.*

Our games $\omega = A$ and B are dominance solvable. Thus, applying the result of Milgrom and Roberts (1991) (see also Nachbar, 1990), we can conclude that whatever the starting point (the actions in the first round) one must converge to the unique Nash equilibrium in the two games.

2) *Convergence in Coarse learning model.*

The convergence result follows a logic similar to that in Milgrom and Roberts (1991) and is the subject of Jehiel (in preparation).

Since, action e is strictly dominated for Column players in both games, it follows that for $t > 5$, there is no e in $x_2^C(t)$. From such a t on, Row players do not play γ in either game since action γ for Row player is strictly dominated in both games once action e is eliminated. Thus from $t > 10$ on, there is no γ in $x_1^F(t)$, and column players do not play actions d and b in A (given that there is no γ in $x_1^F(t)$, d is dominated by a and b is dominated by any strictly convex combination of a and c) and they do not play actions a and c in B (given that there is no γ in $x_1^F(t)$, a is dominated by a strictly convex combination of b and d and c is dominated by b). It follows that from $t > 15$ on, $x_2^C(t)$ contains an equal mixture of a convex combination of a and c (that accounts for Column players' behaviors in A) and a convex combination of b and d (that accounts for Column players' behaviors in B). From $t > 15$ on, row players play action β in A and action α in B (since the other action is dominated in the respective games – this follows from the observation that for all $\lambda \in (0, 1)$, $\frac{1}{2}(15\lambda + 10(1 - \lambda)) - \frac{5}{2} > 0$). From $t > 20$ on, Column players only see action β in A and thus they play a . They only see action α in B and thus they play b . The ABEE is being played from then on. \square

4. Experimental design

The computerized experiments were conducted at the UCL-ELSE Economics Laboratory between February and December 2005, with the sessions for the Solo treatment running in January 2009. Upon arrival at the lab, subjects sat down at a computer terminal to start the experiment. Instructions were presented on the computer screen and a written summary of the instructions was also handed out. Subjects were invited to raise their hands at any time to ask questions which would be answered privately. Subjects were not allowed to take any notes during the entire experiment.

The experiment consisted of five treatments which varied in the accessibility of the information available to subjects about the actions of others. Each session involved eight subjects and four sessions were run for each treatment. In total 160 subjects participated in the experiment, drawn from the student population at UCL. Their subjects of study included a cross section of arts, humanities, social science, science and medical subjects. Subjects were paid a turn-up fee of £5 and in addition to this were given £0.05 per point won during the experiment. The average payment was around £13 per subject, including the turn-up fee. All of the sessions lasted between 45 minutes and 1 hour, with the HARD treatments taking the longest. And subjects took longer to consider their choices at the start of the experiment: generally, over our sessions of 60 rounds, the first 20 rounds took a similar length of time as the last 40 rounds.

In all treatments, subjects were split up equally into two roles, Row and Column. Each session consisted of sixty rounds where Row and Column subjects were randomly matched into four pairs to make a choice in one of two normal form games, the Row subject choosing the row in the game matrix and the Column subject choosing the column. The two normal form games chosen were detailed in Table 1. In all treatments except SOLO, two pairs in each round were allocated to “game A ” and two to “game B ”, and both subjects in each pair knew which game they were in.⁸ In the SOLO treatment, one of the

⁸ Subjects also knew that both games were played with the same frequency.

two games was allocated for all rounds and all subjects for each session (“game A” for two sessions and “game B” for two sessions).

In all treatments, subjects could only see their own payments in each game, and were given information about the choices made by the subjects in the other role in the previous five rounds.

In all treatments the Column subjects were presented in every round with the number of times each row had been chosen, in the current game, over the last five rounds. The number was shown against the row on their payoff matrix, and the experiment instructions explained the meaning of the numbers and that they were being provided “to help you make your decision”. Column subjects were never given any feedback about play in the game not currently seen. For example, if a Column subject was in game A in round 25, she only saw the distribution of choices for game A on the screen from rounds 20 to 24. In a later round, she may have been in game B, and only feedback for game B would have been seen.

The feedback provided to Row subjects varied across treatments. The feedback screens used in FEASIBLE and HARD were shown in Section 2 (see Figs. 1 and 2). Note that the ordering of the grid showing the actions chosen in each of the four matches of each of the last five rounds (see Figs. 1(a) and 2) was randomized independently each round, and subjects were informed of this.⁹ Note also that Row subjects were allowed to consider each screen for as long as they wanted, but once they had clicked to move on to the next screen they could not go back to the previous screen and they were not allowed to take any notes.

The treatments COARSE, FINE and SOLO were simpler, as previously described. In FINE and SOLO, Row subjects were given similar information to that of the Column subjects, i.e., the number of times each column had been chosen over the last five rounds in the game to be currently played. In COARSE the Row subjects saw the total number of column choices aggregated over the two games A and B, and they were not given any information about the game in which the choices were made. In the treatment SOLO the same game was seen by all players in every round for each session.

Before we move on to analyze the data from this experiment, let us briefly discuss the methodological issue of varying feedback information and its presentation or, more generally, interface design. Many economic experiments try to test equilibrium predictions by giving these predictions their “best shot”. Accordingly, instructions, interfaces, and feedback are then designed to present the necessary information in the easiest accessible way such that, if equilibrium predictions are falsified, this points at deep and fundamental biases in decision making and not simply at bad experimental design and confusion. In such experiments, one would never present feedback in the way we do here. However, this is not the only class of economic experiments. Experiments that try to dig deeper into subjects’ decision making processes, learning and heuristics have traditionally exploited variations in feedback design or accessibility of information. Mouselab studies that investigate the extent of backward induction (see Costa-Gomes et al., 2001) or the subjects’ desire to make use of particular feedback (Bigoni, 2010) do precisely that. By providing “obstacles” and observing whether and how subjects try to overcome these, one can gain important insights into subjects’ reasoning processes. Similarly, the provision of different and sometimes seemingly irrelevant types of (feedback) information has proved to be useful to understand how subjects learn in economic environments (see, for example, Huck et al., 1999). Our methodology of making it harder and harder to process objectively identical information in order to understand how subjects adjust their processing of this information falls precisely into that category of scientific endeavor.

5. Results

With the experimental design in mind, we rephrase our research agenda by the following questions: (i) Do the observed behaviors in FINE, COARSE, SOLO, FEASIBLE and HARD stabilize? (ii) Do the observed behaviors in FINE, COARSE, SOLO, FEASIBLE and HARD resemble each other? (iii) Do the long run behaviors in FEASIBLE and HARD resemble the long run behaviors in FINE, COARSE or SOLO? (iv) Do the observed patterns of play in FINE, COARSE, SOLO, FEASIBLE and HARD relate to the Nash equilibrium, the analogy-based expectation equilibrium or none of these?

We proceed by obtaining a positive answer to question 1. We also show that FEASIBLE, FINE and SOLO all look similar, and that HARD looks like COARSE. We go on to provide evidence that behaviors in FEASIBLE, FINE and SOLO are best explained by the Nash equilibrium and behaviors in HARD and COARSE are best explained by the analogy-based expectation equilibrium (in which Column players use the fine analogy grouping and Row players use the coarse analogy grouping, see Section 3).

We split the analysis of the results into two parts. The first considers the aggregate observed behavior in the last half of the experiment, and treats this as equilibrium play, estimating parameters of an ‘analogy-based quantal response equilibrium’ (see Appendix A). The second considers all rounds and looks at the equivalent learning model to allow us to look closer at individual behavior.

5.1. Aggregate data

A first set of summary statistics is given in Table 2. This shows for all five treatments the frequencies of each choice in each of the two games for Row and Column players in the second half of each session. While this table shows averages

⁹ That is, the action played in any of the four matches in any of the five last rounds could appear in any position of the 5×4 matrix of Fig. 1 in HARD and Fig. 2 in FEASIBLE with an equal probability.

Table 2

Summary of choice frequencies for second half of all sessions.

Row player game A	FINE	FEASIBLE	HARD	COARSE	SOLO
α (Nash equilibrium)	0.79	0.88	0.25	0.24	0.73
β (Coarse ABEE)	0.18	0.13	0.66	0.70	0.20
γ	0.04	0.00	0.09	0.06	0.07

Row player game B	FINE	FEASIBLE	HARD	COARSE	SOLO
α (Coarse ABEE)	0.16	0.21	0.69	0.78	0.16
β (Nash equilibrium)	0.80	0.76	0.20	0.14	0.78
γ	0.03	0.03	0.11	0.09	0.06

Col player game A	FINE	FEASIBLE	HARD	COARSE	SOLO
<i>a</i> (Coarse ABEE)	0.41	0.31	0.73	0.83	0.40
<i>b</i>	0.00	0.01	0.00	0.00	0.02
<i>c</i> (Nash equilibrium)	0.50	0.67	0.06	0.10	0.45
<i>d</i>	0.08	0.01	0.22	0.08	0.11
<i>e</i>	0.00	0.00	0.00	0.00	0.01

Col player game B	FINE	FEASIBLE	HARD	COARSE	SOLO
<i>a</i>	0.01	0.00	0.00	0.00	0.00
<i>b</i> (Coarse ABEE)	0.12	0.28	0.48	0.57	0.10
<i>c</i>	0.08	0.05	0.25	0.16	0.00
<i>d</i> (Nash equilibrium)	0.80	0.67	0.27	0.27	0.90
<i>e</i>	0.00	0.00	0.00	0.00	0.00

across all periods it is important to note that after an initial learning stage, the choices of Row and Column players stabilize in the later phases of the experiment. To test this, we considered the row player's choice frequency in each phase of 15 rounds, for each combination of game and treatment. The last quarter of the experiment was compared with the 3 preceding quarters using a Fisher's exact test on the resulting contingency table, at the 1% level. In most cases there was a significant difference between the choice distribution in the first 15 rounds and the last 15 rounds.¹⁰ There was only a significant difference between rounds 31 to 45 and 46 to 60 in two cases – column players in game B in the HARD and SOLO treatments. In these two cases, the play stabilized in the last 15 rounds. We conclude, therefore, that play does stabilize in each of our four treatments.

A first observation we can make by inspecting the table is that there is a significant difference of behaviors across treatments. Second the modal behaviors of Row players and Column players in SOLO, FINE and FEASIBLE coincide with the behaviors arising in the Nash equilibrium. And the modal behaviors of Row players and Column players arising in COARSE and HARD coincide with the ABEE with coarse grouping. The finding that SOLO is similar to FINE supports very directly the intuition that there is no significant learning spillover at work in FINE. The finding that the distributions of behaviors in FINE and FEASIBLE on the one hand and COARSE and HARD on the other are similar, is suggestive that the feedback about opponents' play is accessible game by game in FEASIBLE but only in aggregate over the two games in HARD.

Notice though that there are still some systematic deviations from equilibrium play (either Nash equilibrium for SOLO, FINE and FEASIBLE or ABEE for COARSE and HARD), such as column players choosing *a* in game A in the FINE, FEASIBLE and SOLO treatments. However, in all cases these deviations are the same in FINE and FEASIBLE on the one hand and COARSE and HARD on the other. That the SOLO treatment also looks similar to FINE and FEASIBLE provides evidence that the deviations from equilibrium play in the latter treatments are not due to the mixed games presentation.

To examine these frequencies in more detail, and in particular to ask why some deviations from equilibrium play are more pronounced than others, we apply the analogy-based quantal response equilibrium model detailed in Appendix A. In summary, this consists of a small extension of the quantal response equilibrium concept of McKelvey and Palfrey (1995) to combine it with the analogy based expectation model. Each player is endowed with a noise parameter λ_i which measures the extent to which that player best responds to the expected utilities for each choice. More precisely, if action a_i is perceived to provide an expected utility $\bar{u}_i(a_i)$, action a_i is played with a probability proportional to $\exp \lambda_i \bar{u}_i(a_i)$.

If $\lambda_i = 0$ then the player will equally mix all actions, and the strategy of the player tends towards a best response as $\lambda_i \rightarrow \infty$. The analogy-based quantal response equilibrium is thus defined as strategy profiles which are mutual analogy-based best responses given the noisy best response function just defined and further detailed in Appendix A. Therefore, each pair of λ_i for each player (one for row players and one for column players) allows us to calculate corresponding analogy-based quantal response equilibrium strategy profiles. Our estimation proceeds by maximum likelihood methods, by choosing the values for $\lambda = (\lambda_r, \lambda_c)$ (one for Row players and one for Column players in the five treatments) for which

¹⁰ The exceptions for row players were in game A in the HARD treatment and in game B in the COARSE treatment. For column players the first 15 and last 15 rounds were not significantly different for game B in the FINE, COARSE and FEASIBLE treatments.

Table 3
Estimates of QRE parameters for each model.

Treatment	Fine		Coarse		Vuong
	Row	Col	Row	Col	
FINE	1.846 (0.2078)	1.290 (0.08406)	-3.816 (0.2117)	0.2789 (0.02262)	-4.272
FEASIBLE	2.309 (0.2266)	1.239 (0.07035)	-4.101 (0.2183)	0.2638 (0.02073)	-17.62
HARD	0.1529 (0.02045)	2.717 (0.2344)	1.233 (0.1155)	1.856 (0.1431)	2.364
COARSE	0.2731 (0.03985)	2.201 (0.2273)	1.585 (0.1213)	1.926 (0.1265)	5.454
SOLO	1.260 (0.1336)	1.646 (0.1190)	-	-	-

the probability density of the observed choice frequencies is maximized (with respect to the multinomial distribution of choices implied by the corresponding equilibrium strategies). That is, if the observed frequencies of rows $i = 1, 2, 3$ are n_i and the observed frequencies of columns $j = 1, \dots, 5$ are m_j , and if the equilibrium strategies in the analogy-based quantal response equilibrium for given λ parameters specify that row i is played with probability $q_i(\lambda_r, \lambda_c)$ and column j is played with probability $r_j(\lambda_r, \lambda_c)$, the likelihood function for this one observation $(n, m) = (n_i, m_j)_{i,j}$ is proportional to:

$$L(\lambda_r, \lambda_c | n, m) = \prod_i q_i(\lambda)^{n_i} \prod_j r_j(\lambda)^{m_j} \tag{4}$$

For each treatment, and for both the Fine and Coarse ABEE models we estimate the λ parameters that maximize L given the observations (n, m) . We then apply a test of which model better fits the data. As the two models are not nested, it isn't possible to use a simple likelihood ratio test to compare the goodness-of-fit. We chose to use the goodness of fit test introduced in Vuong (1989). This test compares the differences in the maximal log likelihoods LL_k^f and LL_k^c for the fine and coarse models for each observation (experiment session) k and provides a test statistic asymptotically distributed $N(0, 1)$ under the null hypothesis that both models fit the data equally well. The statistic is evaluated using the following formula (where K denotes the total number of experiment sessions):

$$\frac{\sum_k (LL_k^c - LL_k^f)}{\sqrt{\sum_k (LL_k^c - LL_k^f)^2 - \frac{1}{K} (\sum_k LL_k^c - LL_k^f)^2}} \tag{5}$$

A negative value suggests that the fine feedback equilibrium fits better and a very positive number points towards the coarse feedback being used. The results of this estimation are detailed in Table 3.

Table 3 provides further positive evidence that players are able to use the fine beliefs in the FEASIBLE treatment and the coarse beliefs in the HARD treatment. It's not possible to say that the fine QRE parameters are the same between the FINE, FEASIBLE and SOLO treatments, nor that the coarse QRE parameters are the same between the COARSE and HARD. However, this may not be surprising given the difference between the presentation of the feedback between these treatments.¹¹ What the results do strongly suggest however, is that the fine feedback is being used in the FINE, FEASIBLE and SOLO treatments (perhaps to a greater or lesser extent) and that the coarse feedback is not being used. Similarly, the estimates for the COARSE and HARD treatments are similarly convincing that the coarse feedback only is used in those treatments.

These statistics strongly confirm our previous conclusion that the fine feedback is being used in the FINE and FEASIBLE treatments and that the coarse feedback is being used in the COARSE and HARD treatments. In addition to this, the QRE approach allows us to explain the particularly distinct deviation of column players. Appendix B (Table 7) shows the choice frequencies for the analogy-based QRE for the estimated values of λ for comparison with Table 2.

Overall, the aggregate frequencies provide evidence for the affirmations that i) FEASIBLE and FINE are best explained by Nash equilibrium; ii) HARD and COARSE are best explained by the analogy-based expectation equilibrium (ABEE). The results also strongly support the hypothesis that the accessibility of information determines how the information is used when making decisions. The FEASIBLE treatment shows that when the information on which game corresponds to each choice is more easily accessible, it is used (i.e., separate distributions of opponents' actions are considered for each game, similarly to what happens in the FINE treatment, and learning spillovers do not occur). The HARD treatment makes the specific information on which game corresponds to each action considerably less accessible. But crucially subjects do not throw away this information however coarse it may be. Rather they distill a statistic from that information that coincides with exogenously coarse information. Consequently, on an aggregate level subjects appear to behave in HARD somehow similarly to subjects in COARSE, and these behaviors turn out to be well explained by the analogy-based expectation equilibrium.

¹¹ Also, the equilibrium strategies depend strongly on the λ parameters for both roles, and as such we think that comparing the individual responses to feedback are a better test of the similarity between behavior across the treatments.

Table 4

Best response frequencies for row players.

	None	Coarse	Fine	Both ^a
FINE	0.13	0.17	0.76	0.07
FEASIBLE	0.05	0.21	0.80	0.06
HARD	0.22	0.60	0.43	0.26
COARSE	0.15	0.68	0.28	0.11
SOLO	0.35	–	0.65	–

^a The Fine and Both columns double count those cases where a choice was a best response to both the feedback types. The Both column is included to show the extent to which the incentives overlap for each treatment.

5.2. Individual data

In this subsection, we wish to get a better understanding of the dynamics of actions. Since our main interest lies in the understanding of Row players, we consider the individual choices of the Row subjects. Our objective is to study whether at the individual level, the dynamics of Row players' actions are well described by the learning dynamics with fine or coarse empirical distributions of Columns over the last five rounds as described in Section 3. We first consider the case in which exact best responses are considered. We next move on to analyze the analogous learning models with noisy best-responses. In this class of noisy best-responses, we first treat populations of Row players as if they were homogenous. Finally, we analyze individual subjects separately.

Firstly, we examine whether individual decisions are best responses to the information provided. In doing so we shall distinguish between fine and coarse information, taking into account that in some instances best replies to both types of information may coincide. We consider two different kinds of beliefs that Row players might hold about the strategy of the Column players. These are constructed from the empirical distributions of past play over the previous five rounds. Fine beliefs use the fine feedback, thereby considering only previous play of the Column players for the relevant game. Coarse beliefs rely on the total frequencies of past play of the Column players over both games. These beliefs have been defined formally at the end of Section 3 (see expressions (2) and (3), respectively).

Having constructed these beliefs we are able to calculate the expected payoff for each row. Table 4 details the frequencies of best responses, i.e. instances of row choices that maximize expected payoffs given some method of forming beliefs. The table is very suggestive. Both, in FINE and FEASIBLE, Row subjects best respond to fine information around 70% of the time while in COARSE and HARD Row subjects best respond to coarse information close to 65% of the time.

While the above is highly suggestive, we now refine the analysis by taking care of the varying incentives to best respond for each choice. Perhaps the deviations were more often in situations where little was to be lost from taking another choice? To examine this quantitatively we estimate a fixed-effects logit model for the expected utilities described above. Given the parameters λ^F and λ^C , Row player is viewed as choosing action a_i in game ω with a probability proportional to

$$\exp[\lambda^F \tilde{u}_i(a_i, x_j^F; \omega) + \lambda^C \tilde{u}_i(a_i, x_j^C; \omega)]$$

where $\tilde{u}_i(a_i, x_j; \omega)$ denotes i 's expected payoff in game ω when choosing a_i and j is expected to play according to x_j , and x_j^F and x_j^C represent the expectations as derived from the empirical distributions in the fine and coarse representation, respectively (see Eqs. (2) and (3) in Section 3). For each session, we estimate λ^F and λ^C so as to maximize the likelihood function given the observed choices of Rows over the sixty periods $t = 1, \dots, 60$. We cluster the standard errors at the subject level¹² and the estimated parameters are shown in Table 5.

These results act to confirm the conclusion of the previous section: FEASIBLE looks like FINE and SOLO, and HARD looks like COARSE. However, we can now use Wald tests to say that the parameters for the Nash beliefs in the FEASIBLE and SOLO treatments are not significantly different from those in the FINE treatment, and also that the parameters for the ABEE beliefs are not significantly different between the COARSE and HARD treatments. In addition to this, we can now observe that the parameters for the Nash beliefs in the HARD and COARSE treatments, and also those for the ABEE beliefs in the FINE and FEASIBLE treatments are not significantly different from zero. This adds weight to our conjecture that only the fine feedback is used by subjects in the FINE and FEASIBLE treatments and that only the coarse feedback is used in the HARD and COARSE treatments.

Finally, we look at each subject separately. We repeated the fixed-effects logit regression, but for each subject individually. For each subject we thus obtain an estimate of the λ parameters for the learning models using both fine and coarse feedback. In each case, we tested the hypothesis that the parameters were positive at the 5% level. If a subject had only a significant parameter for the fine feedback model, (s)he was allocated to the type 'Fine' and similarly those whose regression coefficient for the coarse model was the only significant estimate were allocated to the 'Coarse' type. There were a few subjects where both parameters were significant. Here we indicate which parameter was the largest. The frequencies of each type are indicated in Table 6. They provide evidence that some agents are able to distinguish the feedback

¹² In doing so we are assuming that any session level fixed effects are transmitted across subjects solely through the feedback over previous choices.

Table 5

Noise parameters for learning model.

Row	Nash		ABEE	
	Estimate	Std. error	Estimate	Std. error
FINE	1.764*	0.3340 (0)	0.00 ^a	–
FEASIBLE	2.055*	0.1674 (0)	0.1247	0.3044 (0.682)
HARD	0.3609	0.3263 (0.269)	1.093*	0.3835 (0)
COARSE	0.1974	0.3246 (0.499)	1.635*	0.1272 (0)
SOLO	1.320*	0.2766 (0)	–	–
Column (all treatments)	1.336*	0.09397 (0)	–	–

^a Note that the parameter values are constrained to be positive, and in this case the constraint is binding.

Table 6

Player types based on significant parameters in logit estimation.

Treatment	Neither	Fine	Coarse	Both (Fine)	Both (Coarse)
FINE	0	12	2	1	1
FEASIBLE	0	13	0	2	1
HARD	2	4	10	0	0
COARSE	1	2	12	1	0

in the HARD treatment.¹³ We present these results as indicative, but further research is required to draw more definite conclusions.

We also examine briefly the possibility that subjects transform feedback into beliefs in different ways. We consider the class of models where the vector of relative frequencies of opponents' actions that has been observed $f = (f_i)_i$ is mapped into a belief vector $b = (b_i)_i$ through some weakly monotone function. This would allow for interesting "hybrid models". Take for example, a function that assigns 0 to all actions i for which $f_i < \varepsilon$ for some small ε and a belief of $b_i = \frac{1}{k}$ where k is the number of strategies with a relative frequency above ε . This would be the equivalent of a level-1 algorithm where subjects essentially ignore actions that are played rarely.¹⁴ As it happens such a model would fit our aggregate data almost as well as the standard ABEE model (the equilibrium prediction is the same). However, it fails to explain as well the times when the equilibrium action is not chosen (that is, we observe evidence for strict monotonicity of beliefs with respect to the feedback¹⁵).

Overall, the results provide positive responses to all of our research questions. They suggest that players use the fine feedback when it is accessible and, more significantly, they use the coarse feedback when only that is accessible (rather than choosing other ways to play the game such as level-1 reasoning, maxmin or other 'hybrid' approaches). This provides the first positive experimental test of the analogy-based expectation equilibrium in the laboratory.

6. Related literature

There are almost no experimental papers dealing with multiple games at the same time. The only exception we are aware of are the papers by Cooper and Kagel (2007, 2008) which focus on cross-game learning, the ability of subjects to take what has been learned in one game and transfer it to related games.¹⁶ Despite the common feature that several games played by the same subjects are being analyzed experimentally, the two lines of research are different and rather complementary. Cooper and Kagel let subjects learn in one game and then study whether there is transference of learning to a new game

¹³ There are two subjects in the HARD treatment who best respond to the fine feedback more than 80% of the time. Between them they account for most of the choices that were best responses to fine beliefs and not to coarse beliefs.

¹⁴ In the context of our experiment, it is only really possible for subjects to know that an action hasn't been taken in a particular game (when it is absent from the combined feedback). It is not generally possible for subjects to know that an action has been taken in a particular situation (as the relative frequencies of each action rarely exceed 50%).

¹⁵ We compared the expected utilities (derived using the ABEE beliefs) of the action taken and the best alternative for players identified above as being of type Coarse. Using the Mann–Whitney U test, we were able to show that the equilibrium action was chosen more often when the incentives to do so were stronger (or the incentives not to do so were weaker).

¹⁶ They consider limit pricing games such as those analyzed in Milgrom and Roberts (1982), and they study the ability of subjects to learn to play strategically (i.e. engage in limit pricing) for a new specification of the parameters of the model when they have learned to play the equilibrium for another specification of the parameters.

played after the first. The purpose of our experiment is to understand the consequences of feedback spillover, and in particular how it affects long run behaviors in the various games that are being played contemporaneously. Accordingly, unlike in Cooper and Kagel, it is not the case in our experiment that one game is being played (many times) first before one moves to the next game. And the question we address is whether an analogy-based expectation equilibrium is being played in the case of feedback spillover rather than whether if Nash equilibrium is played in one game it increases the chance that Nash equilibrium is played in another game. However, there are similarities in the findings of these different research strands. Cooper and Kagel show how experience in one game can generate strategic sophistication, that is, an immediate better grasp of the strategic implications of changes in the game environment. With longer learning (in the first game), more subjects become sophisticated (in the second game). In our study, we show how subjects adapt their learning heuristics to the difficulty of processing information (their mode of behavior is different in FEASIBLE and HARD). So, in both cases, the presence of multiple games interact with some other cognitive component resulting in different outcomes as one makes the environment more or less complex.

In a related vein, Rick and Weber (2010) demonstrate how, in the absence of feedback, subjects can acquire deeper concepts (here elimination of dominated strategies) that are again transferable from one game to the next. As in Cooper and Kagel and our study, such a finding is suggestive of a cognitive process that is not present in standard models of learning. In the context of an experiment where subjects play different normal form games, Haruvy and Stahl (2009) suggest how EWA learning can sometimes be augmented through re-labeling strategies according to properties they share (from a level- k reasoning perspective) to explain the data. As an alternative approach they also study “rule learning” (Stahl, 1996) whereby subjects acquire new modes of reasoning. These studies like ours should be viewed as suggesting new directions in which to improve learning models.

Given the learning interpretation of the ABEE concept (as made explicit in Section 3 and put to experimental test in Section 5), it is legitimate to place our work into the broader perspective of the experimental literature concerned with learning. An important question addressed by this literature concerns the issue whether reinforcement learning (in which the choice of actions is based on the subject's own past performance) or belief-based learning (in which the choice of actions is based on the opponents' behaviors) fit best the experimental data (see Erev and Roth, 1998; Camerer and Ho, 1999; Wilcox, 2006).¹⁷ In our experiment subjects are not informed of their past performance (until the very end), hence we abstract from these issues and force the relevant learning models to belong to the family of belief-based learning models.

Our investigation has also some connection with the experimental literature interested in the number of reasoning steps made by subjects in non-repeated interactions (see Stahl, 1993; Nagel, 1995; Costa-Gomes et al., 2001; Camerer et al., 2004). Key differences with our approach are that we are interested in long run behaviors in which there has been plenty of time for behavior adjustment, and there is no idea of feedback spillover in the level- k type of reasoning. In connection to these approaches, it should be mentioned that the introspective reasoning steps considered in the level- k literature require (except for levels 0 and 1) that players know the payoffs of their opponents, which as already mentioned is not the case in our experiment.

Finally, our experiment has also some remote connection with those experiments that try to understand which games players perceive to be playing.¹⁸ In our experiment though, the misperception of Row subjects observed in COARSE and HARD concern the play of Column subjects rather than of the game they are playing. As such, the effect of feedback spillover as experimentally explored in our paper is not of viewing our subjects as playing a Nash equilibrium of a different game but rather viewing the players as playing a different notion of equilibrium of the original (set of) games.

7. Conclusion

This paper has provided a first positive test for the analogy-based expectation equilibrium. Yet, it should be viewed as a first step and more work should be devoted to the more general enterprise of understanding how agents process data in environments with more than one task. Such data sets are generally huge and often unstructured – much in contrast to how feedback is normally structured in experiments.

We leave for future research the task of exploring how our results related to feedback spillover would be affected if subjects received information about their own performance during the experiment. It would certainly also be interesting to explore how long run behaviors are affected by feedback spillover when players are informed of the payoffs of their opponents. This would open the door to the possibility that upon receiving the feedback subjects make further inferences based on introspective reasoning about the opponent's incentives (see Ehrblatt et al., 2005, for an explicit account of this).

It might also be nice to explore how subjects categorize data from different games in order to form their expectation. Do subjects somehow have a sense of which games are more similar to one another, or are feedback spillovers the outcome of some external force (as was the case in our experiment – the external force being the experimentalists)? Presumably, both forms of feedback spillover (individually chosen or externally induced) are at work in the real world. The present study which has put the emphasis on the external component of feedback spillover opens the door to the possibility that a

¹⁷ We note that while the early contributions have suggested that reinforcement learning explains better the data, this has been recently challenged in Wilcox (2006).

¹⁸ See Devetag and Warglien (2008) for a recent experiment on this (and also Oechssler and Schipper, 2003, for another perspective).

designer interested in obtaining some desirable outcomes may use feedback spillover as an instrument (which is typically not considered in mechanism design). Alternatively, if agents make their decisions based on feedback spillover (as resulting from individual choices), then a designer should take this into account as it may affect the performance of the mechanism she selects. More theoretical and experimental work along these lines is needed.

As we briefly considered earlier, one could also examine an extension of the ABEE concept where subjects use other heuristics to transform feedback into beliefs in weakly monotonic ways. These “hybrid models” and the situations in which they are tested would have to be designed carefully to allow them to be distinguished from the baseline ABEE model and other existing learning models from the literature. Such models could fit equilibrium predictions almost equally well as the standard ABEE model but would differ in their predictions for the dynamics that we observe before play settles down and, as in our experiment, in how and where players deviate from optimal play. We see no evidence in our data to support these hybrid models but one could imagine that such models could prove useful in more complex environments.

Appendix A

In this appendix, we extend the definitions developed in Section 3 to allow players to use quantal best-responses rather than ordinary best-responses. We adopt the same notation as in Section 3, and we further define:

$$\bar{u}_i(a_i; \omega) = \sum_{a_j \in A_j} \bar{\sigma}_j(a_j | \omega) u_i(a_i, a_j; \omega) \tag{A.1}$$

where $\bar{\sigma}_j$ is as defined in (1).

Definition. A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is an analogy-based quantal response equilibrium given the analogy partitions An_1, An_2 and the logit parameters λ_1, λ_2 if for all $i, \omega \in \Omega$ and $a_i \in A_i$,

$$\sigma_i(a_i | \omega) = \frac{\exp \lambda_i \bar{u}_i(a_i; \omega)}{\sum_{a'_i} \exp \lambda_i \bar{u}_i(a'_i; \omega)}$$

where $\bar{u}_i(a'_i; \omega)$ is given by (A.1).

The learning models defined in Section 3 are extended analogously to cope with quantal best-responses rather than ordinary best-responses. That is, for every $x_j \in \Delta A_j$, let

$$\tilde{u}_i(a_i, x_j; \omega) = \sum_{a_j} x_j(a_j) u_i(a_i, a_j; \omega)$$

where $x_j(a_j)$ denotes the weight assigned to a_j in x_j . Define further the probabilities:

$$QBR_i^\omega(x_j | \lambda_i)(a_i) = \frac{\exp \lambda_i \tilde{u}_i(a_i, x_j; \omega)}{\sum_{a'_i} \exp \lambda_i \tilde{u}_i(a'_i, x_j; \omega)}$$

In round $t + 1$, player i chooses action a_i with probability $QBR_i^\omega(x_j | \lambda_i)(a_i)$ where x_j is as defined in (2) or in (3) depending on whether the fine or the coarse feedback case is considered.

Appendix B

Table 7
Expected choice frequencies for the estimated analogy-based QRE.

Row player game A	FINE (fine)	FEASIBLE (fine)	HARD (coarse)	COARSE (coarse)	SOLO (fine)
α (Nash equilibrium)	0.68	0.75	0.27	0.23	0.63
β (Coarse ABEE)	0.27	0.22	0.55	0.62	0.28
γ	0.05	0.03	0.18	0.15	0.09

Row player game B	FINE (fine)	FEASIBLE (fine)	HARD (coarse)	COARSE (coarse)	SOLO (fine)
α (Coarse ABEE)	0.24	0.19	0.62	0.70	0.26
β (Nash equilibrium)	0.72	0.78	0.17	0.12	0.66
γ	0.04	0.02	0.21	0.17	0.08

(continued on next page)

Table 7 (continued)

Col player game A	FINE (fine)	FEASIBLE (fine)	HARD (coarse)	COARSE (coarse)	SOLO (fine)
a (Coarse ABEE)	0.28	0.23	0.77	0.87	0.30
b	0.10	0.10	0.04	0.02	0.07
c (Nash equilibrium)	0.58	0.64	0.10	0.06	0.60
d	0.02	0.02	0.07	0.04	0.02
e	0.02	0.02	0.01	0.01	0.01

Col player game B	FINE (fine)	FEASIBLE (fine)	HARD (coarse)	COARSE (coarse)	SOLO (fine)
a	0.08	0.08	0.05	0.04	0.06
b (Coarse ABEE)	0.16	0.14	0.47	0.58	0.14
c	0.02	0.02	0.18	0.12	0.01
d (Nash equilibrium)	0.73	0.75	0.28	0.24	0.78
e	0.01	0.01	0.03	0.02	0.01

References

- Battigalli, P., 1987. Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali. Unpublished undergraduate dissertation, Bocconi University, Milano.
- Bigoni, M., 2010. What do you want to know? Information acquisition and learning in experimental Cournot games. *Res. Econ.* 64 (1), 1–17.
- Camerer, C., Ho, T.H., 1999. Experience-weighted attraction learning in normal form games. *Econometrica* 67, 827–874.
- Camerer, C., Ho, T.H., Chong, J.K., 2004. A cognitive hierarchy model of games. *Quart. J. Econ.* 119 (3), 861–898.
- Cooper, D.J., Kagel, J.H., 2007. The role of context and team play in cross-game learning. *J. Europ. Econ. Assoc.* 7 (5), 1101–1139.
- Cooper, D.J., Kagel, J.H., 2008. Learning and transfer in signaling games. *Econ. Theory* 34 (3), 415–439.
- Costa-Gomes, M., Crawford, V., Broseta, B., 2001. Cognition and behavior in normal-form games: An experimental study. *Econometrica* 69, 1193–1235.
- Dekel, E., Fudenberg, D., Levine, D.K., 2004. Learning to play Bayesian games. *Games Econ. Behav.* 46, 282–303.
- Devetag, M.G., Warglien, M., 2008. Playing the wrong game: An experimental analysis of relational complexity and strategic misrepresentation. *Games Econ. Behav.* 62 (2), 364–382.
- Ehrblatt, W.Z., Hyndman, K., Ozbay, E.Y., Schotter, A., 2005. Convergence: An experimental study. Mimeo NYU.
- Erev, I., Roth, A.E., 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *Amer. Econ. Rev.* 88 (4), 848–881.
- Ettinger, D., Jehiel, P., 2010. A theory of deception. *Amer. Econ. J. Microecon.* 2, 1–20.
- Haruvy, E., Stahl, D., 2009. Learning transference between dissimilar symmetric normal-form games. Mimeo.
- Higgins, E.T., 1996. Knowledge activation: Accessibility, applicability, and salience. In: Higgins, E. Tory, Kruglanski, Arie W. (Eds.), *Social Psychology: Handbook of Basic Principles*. Guilford Press, New York, pp. 133–168.
- Huck, S., Normann, H.-T., Oechssler, J., 1999. Learning in Cournot oligopoly: An experiment. *Econ. J.* 109, C80–C95.
- Huck, S., Normann, H.-T., Oechssler, J., 2004. Two are few and four are many. *J. Econ. Behav. Organ.* 53, 435–446.
- Jehiel, P., 2005. Analogy-based expectation equilibrium. *J. Econ. Theory* 123, 81–104.
- Jehiel, P., in preparation. Analogy-based expectation dominance solvability and fictitious play dynamics.
- Jehiel, P., Koessler, F., 2008. Revisiting games of incomplete information with analogy-based expectations. *Games Econ. Behav.* 62 (2), 533–557.
- Khaneman, D., 2003. Maps of bounded rationality: Psychology for behavioral economics. *Amer. Econ. Rev.* 93, 1449–1475.
- Lieberman, B., 1960. Human behavior in a strictly determined 3×3 matrix game. *Behav. Sci.* 5, 317–322.
- McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. *Games Econ. Behav.* 10, 6–38.
- Milgrom, P., Roberts, J., 1982. Predation, reputation, and entry deterrence. *J. Econ. Theory* 27, 280–312.
- Milgrom, P., Roberts, J., 1991. Adaptive and sophisticated learning in normal form games. *Games Econ. Behav.* 3 (1), 82–100.
- Nachbar, J., 1990. 'Evolutionary' selection dynamics in games: Convergence and limit properties. *Int. J. Game Theory* 19, 59–89.
- Nagel, R., 1995. Unraveling in guessing games: An experimental study. *Amer. Econ. Rev.* 85, 1313–1326.
- Oechssler, J., Schipper, B., 2003. Can you guess the game you're playing? *Games Econ. Behav.* 43, 137–152.
- Rick, S., Weber, R.A., 2010. Meaningful learning and transfer of learning in games played repeatedly without feedback. *Games Econ. Behav.* 68 (2), 716–730.
- Stahl, D.O., 1993. Evolution of smartⁿ players. *Games Econ. Behav.* 5, 604–617.
- Stahl, D.O., 1996. Boundedly rational rule learning in a guessing game. *Games Econ. Behav.* 16 (2), 303–330.
- Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. *Science* 211, 453–458.
- Vuong, Quang H., 1989. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57, 307–333.
- Wilcox, N.T., 2006. Theories of learning in games and heterogeneity bias. *Econometrica* 74 (5), 1271–1292.