

Learning Spillover and Analogy-Based Expectations: A Multi-Game Experiment

Steffen Huck* Philippe Jehiel† Tom Rutter‡

October 31, 2007

Abstract

We consider a multi-game interactive learning environment and ask ourselves whether long run behaviors in one game are affected by behaviors in the other, i.e whether there are learning spillovers. Our main finding is that learning spillovers arise whenever the feedback provided to subjects about past play is not easily accessible game by game and thus subjects get a more immediate impression about aggregate distributions. In such a case, long run behaviors stabilize to an analogy-based expectation equilibrium (Jehiel 2005), thereby suggesting how one should broaden the notion of equilibrium to cope with learning spillovers.

KEYWORDS: Analogy-based expectation; information processing; experiments; accessibility; interactive learning.

JEL CLASSIFICATION: C72; D82.

1 Introduction

In real life, agents are typically faced with multiple interactions at the same time. Most models in economics abstract from this complexity by focusing on one interaction at the time, and by studying the equilibrium that emerges from this interaction in isolation. Equilibrium analysis implicitly requires that agents have sufficient knowledge of the precise nature of the interaction, perhaps acquired through learning and experimentation, so that behaviors can stabilize. Yet, when agents are engaged in several interactions at the same time, the assumption that one game can be cleanly separated from the other appears optimistic. In general, it seems plausible that behaviors in one interaction might affect the learning dynamics in another

*UCL & ELSE.

†PSE, UCL & ELSE. Corresponding author at address: PSE, 48 Bd Jourdan, 75014 Paris. email: jehiel@enpc.fr.

‡UCL & ELSE.

and vice versa. If this is the case, the notion of equilibrium should be amended. In other words, if there are learning spillovers between several interactions the usual notion of Nash equilibrium may be inappropriate to describe behavior in the long run and a new notion of equilibrium may be required.

The extent to which learning spillovers are relevant will, intuitively, depend on the information that agents receive in the learning phase. If feedback is such that precise information for each different interaction is readily available and easily separated, spillovers may be minimal. If, on the other hand, information is scattered or not cleanly separated for each game, in short, if precise information is less accessible, agents might rather rely on coarse (heuristic) statistics to inform themselves, which, in turn, may induce spillovers.

In this paper we present an exploratory experimental study into these issues. We consider a multi-game experiment, vary feedback information and its presentation and check under which conditions we observe spillovers. Specifically, we consider an experimental setting in which subjects interact in two different two-player games, *A* and *B*, over 60 periods (each player plays each game 30 times). Both games *A* and *B* have a Column player choosing between five potential columns and a Row player choosing between three potential rows. Subjects were assigned to the role of Row player or Column player for the entire experiment and they were randomly matched in each round of the experiment. Games *A* and *B* were chosen to be simple dominance solvable games so that in the absence of learning spillovers, convergence to Nash equilibrium should be expected.¹

As mentioned above, whether or not learning spillovers occur is likely to depend on the feedback received by subjects in the learning phase. Hence, the structure and presentation of feedback is the central point in our experimental design. In all our treatments, subjects know how their own payoffs depend on strategy combinations, but not how their opponent's payoffs depend on these.² In addition, after each round of play, subjects receive feedback about the previous choices over the last five rounds of the subjects assigned to their opponent's role.³ Column players always receive precise feedback, i.e., feedback that is separated for games *A* and *B*. Hence, for them we do not expect learning spillovers. Our treatments differ only in the feedback given to Row players. In two treatments called FINE and COARSE, information

¹Any reasonable belief-based learning model would predict so. Such a conclusion should be contrasted with the ones arising with games admitting a unique Nash equilibrium in mixed strategies such as the ones studied in Erev and Roth (1998).

²In this way one can be sure that subjects' choices of action are driven by the feedback (and the nature of players' own payoffs) rather than by introspective reasoning about the incentives of the opponent, which is the subject of extensive study elsewhere (e.g., Stahl, 1993, Nagel, 1995, Costa-Gomes et al., 2001, and Camerer et al., 2004). It may be argued that the feedback obtained by subjects may help them assess their opponent's incentives. But, even in this case, the ultimate choice can only be based on their own payoff structure and the feedback they receive.

³We chose not to let subjects be informed about their own performance until the end of the experiment so as to avoid the issue whether subjects rely on belief-based learning-like or reinforcement learning-like procedures which is the subject of extensive study elsewhere (see Erev and Roth, 1998, and Camerer and Ho, 1999).

presentation is ‘ideal’—it could not be presented in a way that makes it easier to digest. In treatment FINE, information is moreover nicely separated for each game. Thus, in FINE, we do not expect any learning spillovers, and we expect convergence to Nash equilibrium. In treatment COARSE, information is again presented very cleanly but this time it is exogenously coarse, that is, *aggregated* across the two games. In other words, in COARSE we enforce learning spillovers (simply because information for each game cannot be separated).

Such ideal information where all relevant statistics appear at precisely the right time and next to the objects they refer to is, while common practice in laboratory experiments, rather rare in real life. This is why we have designed two further treatments in which the same information is offered to subjects as in treatment FINE but in more realistic fashion, i.e., less easily accessible. The two treatments which will be described in much detail below differ in how hard it is to access the relevant information. The treatments’ names, FEASIBLE and HARD, indicate our expectation about the accessibility of precise information.

Our findings are as follows. First, we observe that in all our treatments the pattern of play stabilizes (roughly after 15 rounds). Second, in different treatments it stabilizes at different action profiles. As expected, in treatment FINE, behavior converges to Nash equilibrium. In treatment COARSE, behavior stabilizes at action profiles clearly different from Nash equilibrium, thereby illustrating that a different equilibrium notion may be required in the presence of learning spillovers. What we observe for the more realistic cases of not perfectly presented information is as follows. In treatment FEASIBLE, behavior converges to Nash equilibrium, suggesting that there are no learning spillovers in this case. In treatment HARD, behavior stabilizes at the same action profiles as in COARSE, thereby showing that there are learning spillovers in this case despite the fact that precise information is available.

From a purely experimental viewpoint, the difference of behavior in FEASIBLE and HARD is a manifestation of framing effects, since the feedback in the two treatments is objectively the same in the sense that the same informational content is present in both cases. The similarity between COARSE and HARD suggests that in HARD, Row subjects process information as if they were only provided with the aggregate distributions of behaviors over the two games.

From a conceptual viewpoint, our results are nicely interpreted using Jehiel’s (2005) analogy-based expectation equilibrium theory. In this theory, players bundle the various possible interactions or games into analogy classes and they form expectations only about the *aggregate* behavior of opponents in each such class. Players are assumed to best-respond against these aggregate beliefs in each game, and, in equilibrium, the beliefs coincide with the aggregate play in each class.⁴ As it turns out, the limit action profiles found in COARSE and HARD correspond mostly to the analogy-based expectation equilibrium associated with the coarse analogy structure

⁴This theory is developed in Jehiel (2005), Jehiel and Koessler (2007) and Ettinger and Jehiel (2006).

for Row players and the fine analogy structure for Column players. Besides, considering the frequency with which subjects best-respond to the empirical distribution of actions over the last five rounds for the game they play (fine beliefs) or over both games (coarse beliefs), we find that Row subjects mostly best-respond to the fine beliefs in FINE and FEASIBLE,⁵ and mostly best-respond to the coarse beliefs in COARSE and HARD. However, we also observe a few subjects in HARD who appear to be best-responding to fine beliefs, which reveals some heterogeneity among subjects' ability to process information.⁶

In the next section we focus on the different feedback scenarios that we implemented in our experiment. This is the core of our design and essential for any conjectures about learning spillovers and possible equilibrium outcomes. In Section 3 we then move on to describe Jehiel's analogy-based expectation equilibrium and apply it to the environment considered in our experiment. 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 information structure

As explained in the introduction, treatments FINE and COARSE presented the information feedback in a rather ideal fashion. In both FINE and COARSE, when making their choices Row players were shown the payoff matrix (with their payoffs only) and at the top of each of the five columns some numbers appeared showing how often during the last five rounds Column players had chosen each column. In treatment FINE the numbers concerned only those matches in which the same game as the one to be played now was drawn. In treatment COARSE the numbers aggregated the choices over games *A* and *B*, and subjects were informed of that (see the appropriate tables in the appendix). In both cases, it was extremely easy to form backward-looking beliefs: past frequencies were readily available and shown right next to the strategies they referred to.

In both our main treatments, FEASIBLE and HARD, precise game-specific information was available. However, it was presented differently. In both treatments, each column in the matrix game was given a color, and Row players were shown the distribution over the last five rounds of Column players' actions using colored

⁵Ex post, we conjecture that in FEASIBLE, Row subjects behaved as in FINE because the Nash equilibrium being in pure strategies subjects managed to identify a dominating colour per game.

⁶We note that in line with the literature on ambiguity (Knight (1921), Ellsberg (1961)), one might have conjectured that faced with coarse or hardly accessible feedback, subjects would have discounted more the information provided to them (for example, by adopting a minmax strategy based on the assumption that their opponent wants to minimize their payoffs). We see no evidence of this in our experiment.

boxes containing letters. The color of each box referred to a choice of column and the letters referred to the game (A or B) in which the corresponding column was chosen. The feedback in HARD was 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 Figure 1a, Row subjects could see the distribution of actions (using color boxes) over the last five rounds, i.e., they were shown twenty coloured boxes, each representing one past choice of a column. However, on that screen alone they could not infer the type of game in which a particular action was chosen. This information was presented on the next screen, shown in Figure 1b, where subjects could 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 (see the detailed description in the appendix).⁷ If subjects perfectly memorized the color pattern from screen 1, they could disentangle the distribution for both games from screen 2. If they had difficulties memorizing the precise pattern, they might, 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 (that leads to coarse beliefs) is more easily accessible than the separate distributions are (yielding fine beliefs).⁸

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

As our results summarized below will make clear, Row subjects in HARD seem mostly to memorize the first feedback screen while paying little attention to the second feedback screen even if as mentioned in the introduction there are a few Row subjects who seem to make use of both screens (see the individual data analysis in Section 5). We will demonstrate this by showing the similarity between HARD and COARSE in which Row subjects got only to know the aggregate distribution of Column players' actions over the two games A and B .

By contrast, in FEASIBLE, the feedback was 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 Figure 2). As our results will make clear, subjects managed in FEASIBLE to play as if they had access to the fine distribution of actions by game, and this will be demonstrated by showing the similarity between FEASIBLE and FINE.⁹

⁷The grey boxes in Figure 1a and the question marks in Figure 1b stand for those cases in the first five rounds in which some matches are missing.

⁸See Higgins (1996) or Khaneman (2003) for an exposition of the accessibility idea.

⁹When we designed treatment FEASIBLE, we thought information processing would be hard, and we were surprised to see convergence to Nash equilibrium. Treatment HARD was designed after seeing the results in FEASIBLE.




Figure 1a

ABB?ABAAAB?ABA?BABB?

Figure 1b




Figure 2

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 model that will be used as a theoretical reference point when we analyze the data from our experiment.

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.

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 $a_j \in A_j$

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

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 $p(\omega')$.

Definition A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is an analogy-based expectation equilibrium (ABEE) 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).

In other words, in equilibrium, for each $\omega \in \Omega$, each player i plays a best-response to the beliefs that player j behaves in each game $\omega' \in \alpha_i(\omega)$ according to the aggregate behavior $\bar{\sigma}_j(\omega)$.

In our experiment, we considered two games $\omega = A$ and B whose payoff matrices for the Row ($i = 1$) and Column player ($i = 2$) are depicted below:

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

Table 1: Normal form games used in the experiment

The two games A and B were played with the same frequency so that $p(A) = p(B) = \frac{1}{2}$. Given the feedback provided to Column player (to be described more precisely in the next Section), the relevant analogy partition for Column player was the fine partition $An_2^f = \{\{A\}, \{B\}\}$. For the Row player, two analogy partitions can be considered: either the fine partition, $An_1^f = \{\{A\}, \{B\}\}$, or the coarse partition, $An_1^c = \{\{A, B\}\}$.

When the Row player's analogy partition is fine

When both players use the fine analogy partition, ABEE coincides with Nash equilibrium. In game A the only Nash equilibrium requires that the Row player plays α , $\sigma_1(A) = \alpha$, and the Column player plays c , $\sigma_2(A) = c$. In game B the only Nash equilibrium requires that the Row player plays β , $\sigma_1(B) = \beta$, and the Column player plays d , $\sigma_2(B) = d$. This is easily seen as the iterated elimination of strictly dominated strategies single out a unique pair of actions in each game.¹⁰

When the Row player's analogy partition is coarse

When $An_1 = An_1^c$ and $An_2 = An_2^f$, the only (analogy-based expectation) equilibrium is (σ_1, σ_2) where $\sigma_1(A) = \beta$, $\sigma_1(B) = \alpha$ and $\sigma_2(A) = a$, $\sigma_2(B) = b$. That is, the Row player plays β in game A and α in game B ; the Column player plays a in game A and b in game B . Thus, the strategies are markedly different from the Nash equilibrium strategies.

It is easy to understand why this 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 $p(A) = p(B) = \frac{1}{2}$).

¹⁰In 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) 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.

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 a routine exercise to check that there is no other equilibrium in this case.¹¹

Learning models

Jehiel (2005) (see also Jehiel and Koessler, 2007) motivates the analogy-based expectation approach by a learning story. More precisely, Jehiel (2005) interprets the ABEE as the limiting outcome of a learning process in which each player i would base his strategy on the sole feedback about the aggregate play of player j in the various analogy classes of An_i .

In our experiment, each session of each treatment consisted of several rounds $t = 1, 2, \dots$. In each round, four Column players and four Row players were randomly matched to play one of the games $\omega = A$ or B , and each game was 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 was 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 concerned the behaviors of opponents (i.e., subjects assigned to the role of the other player) in the last five rounds.¹² 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 $m = 1, 2, 3, 4$,

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

¹¹A similar argument to that for the fine equilibrium suffices. Action e is strictly dominated for the Column player in both games. After eliminating e action γ is strictly dominated in both games. Column players have a fine analogy partition so we can again eliminate actions d and b in A and a and c in B . Finally, we use the analogy partitions to note that if the column players choose 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. The equilibrium follows from the column player's best response to this.

¹²In the first five rounds, the feedback covered the entire history of play.

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:

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

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.

In the coarse feedback case, we have:¹³

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

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, Row players follow the fine feedback belief dynamics, and in the other they follow the coarse feedback belief dynamics.

When Row players follow the fine feedback belief dynamics, this corresponds to a fictitious play dynamics with a five period window. Our games $\omega = A$ and B are dominance solvable. Thus, applying Milgrom and Roberts (1991)'s result (see also Nachbar, 1990), we can conclude that whatever the starting point (the actions in the first five rounds) one must converge to the unique Nash equilibrium in the two games.

No general convergence result is available for the learning dynamics in which Row players follow the coarse feedback belief dynamics.¹⁴ Yet, for the payoffs prevailing in games A and B , it can be checked that whatever the initial actions (in the first five rounds and in the four matches) the system converges to the unique ABEE (with coarse grouping for the Row player), which has been described above.

¹³An essential difference with the (stochastic) fictitious play learning model with cross over considered in Cooper and Kagel (2007a) is that the behaviors and beliefs in the two games influence each other at the same time: there is no prior belief attached to one game that can be exported to the other.

¹⁴It is readily verified though that in general if there is convergence it must be to an ABEE.

4 Experimental Design

The computerized experiments¹⁵ were conducted at the UCL-ELSE Economics Laboratory between February and December 2005. 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. Besides, subjects were not allowed to take any notes during the entire experiment.

The experiment consisted of four 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 128 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 each round, two pairs were allocated to “game A” and two to “game B”, and both subjects in each pair knew which game they were in.¹⁶ 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” (see the first table in the appendix). 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

¹⁵The experiments were implemented using bespoke software developed at the ELSE center. The software used is available at <http://www.homepages.ucl.ac.uk/~uctptdr/expecon.zip>

¹⁶Subjects also knew that both games were played with the same frequency.

been seen. Appendix A shows the instruction sheets handed out at the start of the experiment, which also show screenshots from the experiment software.

The feedback screens used in FEASIBLE and HARD were shown in Section 2 (see Figures 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 Figure 1a and Figure 2) was randomized independently each round, and subjects were informed of this.¹⁷ Note also that Row subjects were allowed to consider a given screenshot 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 two treatments COARSE and FINE were simpler, as previously described. In FINE, Row subjects were given similar information to that of the Column subjects, i.e., the number of times each row 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.

5 Results

With this experimental design in mind, we rephrase our research agenda to ask the following questions of our results. Positive answers to these questions are, we suggest, strong evidence for the existence of learning spillovers and their dependence on framing effects.

Question 1 Do the observed distributions of play in treatments FINE, COARSE, FEASIBLE and HARD stabilize?

Question 2 Do the long run behaviors in FINE, COARSE, FEASIBLE and HARD differ from each other?

Question 3 Do the long run behaviors in FEASIBLE and HARD resemble the long run behaviors in FINE or COARSE?

Question 4 Do the observed patterns of play in FINE, COARSE, FEASIBLE and HARD relate to the analogy-based expectation equilibrium?

We proceed by obtaining (qualified) positive answers to Questions 1-2. We also show that FEASIBLE looks like FINE and HARD looks like COARSE, and that behaviors in FEASIBLE and FINE 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

¹⁷That is, the action played in any of the four matches in any of the five last rounds could appear in any position of the 5x4 matrix of figure 2 in FEASIBLE or figure 3 in HARD with an equal probability.

players use the coarse analogy grouping, see Section 3). We first show this on an aggregate level, and then follow the actions of individual participants through our experiment.

5.1 Aggregate data

A first set of summary statistics is given in Table 2. This shows for all four treatments the frequencies of each choice in each of the two games for Row and Column players. While this table shows averages across all periods it is important to note that after an initial learning stage there is no significant difference between the distributions of choices by Row players in the later phases of the experiment.¹⁸ That is, 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 FINE and COARSE coincide with the behaviors arising in the Nash equilibrium and the ABEE with coarse grouping, respectively. Third, the distributions of behaviors in FINE and FEASIBLE on the one hand and COARSE and HARD on the other are similar, suggesting 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, such as column players choosing a in game A in the FINE and FEASIBLE 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.¹⁹

Overall, the aggregate frequencies of actions suggest positive answers to Question 2 and to the affirmations that—i) FEASIBLE and FINE resemble each other and are best explained by Nash equilibrium; ii) HARD and COARSE resemble each other and are best explained by the analogy-based expectation equilibrium (ABEE).

An important issue is, of course, whether the aggregate frequencies are coincidental or can be traced to individual behavior. In order to investigate this question 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 three different kinds of beliefs that Row players might hold about the strategy of the Column players. The first two are constructed from the

¹⁸We 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 chi-squared test on the resulting contingency table, at the 1% level. In all but 2 cases the first 15 rounds were significantly different from the last 15. However, for rounds 16 to 30, there was no significant difference from the final rounds in all but two cases (game B in both FINE and HARD). There was never a significant difference between rounds 31 to 45 and 46 to 60.

¹⁹We think the deviation to a in treatments FINE or FEASIBLE for Column players in game A can be explained by the fact that the incentive to play c vs a is relatively weak.

Row Player game A	FINE	FEASIBLE	HARD	COARSE
α (Nash Equilibrium)	.73	.80	.24	.24
β (Coarse ABEE)	.18	.19	.64	.66
γ	.09	.02	.12	.10

Row Player game B	FINE	FEASIBLE	HARD	COARSE
α (Coarse ABEE)	.16	.23	.62	.73
β (Nash equilibrium)	.76	.73	.21	.17
γ	.08	.04	.17	.10

Col Player game A	FINE	FEASIBLE	HARD	COARSE
a (Coarse ABEE)	.43	.33	.66	.75
b	.00	.01	.00	.00
c (Nash equilibrium)	.47	.60	.08	.12
d	.10	.05	.25	.13
e	.00	.00	.00	.00

Col Player game B	FINE	FEASIBLE	HARD	COARSE
a	.00	.00	.00	.00
b (Coarse ABEE)	.15	.29	.42	.51
c	.08	.07	.30	.16
d (Nash equilibrium)	.77	.64	.28	.33
e	.00	.00	.00	.00

Table 2: Summary of choices made

Row	None	Uniform	Coarse	Fine	Coarse & Fine
FINE	.05	.08	.17	.76	.07
FEASIBLE	.02	.03	.21	.80	.06
HARD	.08	.15	.60	.43	.26
COARSE	.05	.10	.68	.28	.11

Table 3: Best response frequencies

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). To these beliefs we add a third, the uniform belief, which assumes a uniform strategy of the column players. This could also be interpreted as the Row players ignoring the feedback and taking the row with the highest average payoff.²⁰ The uniform belief corresponds to the “Level 1” reasoning of Camerer, Ho and Chong (2004), with Level 0 employing a uniform strategy. Higher-level beliefs in this model would require knowledge of the payoffs of the opponent, which are not available to subjects. Hence, we cannot consider them here. Having constructed these beliefs we are able to calculate the expected payoff for each row. Table 3 details the frequencies of best responses,²¹ 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 80% of the time while in COARSE and HARD Row subjects best respond to coarse information around 65% of the time. It is perhaps not surprising that the numbers for the COARSE treatment are lower than those for the FINE treatment—after all fine information is more reliable than coarse information. Moreover, the number is lowest in HARD, which may be attributed to the fact that in this treatment even accessing the coarse information requires some substantial cognitive effort and is much harder than in the case where the coarse information is simply given exogenously.²²

In a further test of the similarity or difference between the treatments we directly

²⁰Conveniently, the best response to the uniform beliefs is always γ in both games, meaning that near the suggested equilibria the best responses to our three candidates for beliefs span the available action space.

²¹Coarse and Fine were the only beliefs to give coinciding best responses in more than one round for Row players. The Coarse and Fine columns double count these choices, and this column shows how often the two coincide.

²²An alternative explanation might be that because the information is hard to process, subjects tend to discount it more possibly using their maxmin strategy (the maxmin strategy is obtained by solving the corresponding zero-sum game in which the row subject’s payoff matrix is kept unchanged - in both games row player’s maxmin strategy requires playing the Nash equilibrium strategy with probability 1/3, the ABE equilibrium strategy with probability 4/9 and γ with probability 2/9). This will be further discussed later on.

consider the distribution of expected losses with respect to the fine and coarse models of opponents' play. We calculated the theoretical expected losses for the choices taken, if the Row players were to treat the two kinds of beliefs as the true distribution of column play. Hence, for a subject who best replies against a given belief the expected loss of taking this action is zero. For all non-best replies the expected loss is simply the expected payoff from the chosen action (given some beliefs) minus the expected payoff of the best reply (given the same beliefs). Figure 3 shows a histogram with the empirical distribution of expected losses implied by the choices made by the Row players, when using the fine beliefs, i.e. the past distribution of Column players' actions in the current game. It shows a clear similarity between the FINE and FEASIBLE treatments on the one hand, and the COARSE and HARD treatments on the other. The next graph, Figure 4, shows the expected loss of the choice taken when using the coarse beliefs, i.e. the past distribution of Column players' actions over both games *A* and *B*. Again, this clearly shows the similarity between the FINE and FEASIBLE treatments on the one hand, and the COARSE and HARD treatments on the other.

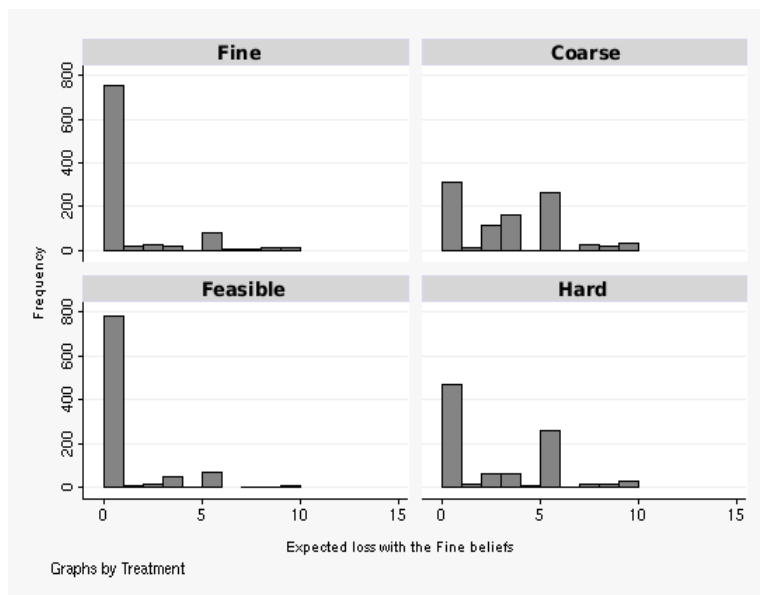


Figure 3: Expected loss using fine beliefs

The results 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.




Figure 4: Expected loss using coarse beliefs

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 very similarly to subjects in COARSE, and these behaviors turn out to be well explained by the analogy-based expectation equilibrium.

5.2 Individual data

Next we consider individual behavior of the Row subjects. Our objective is to see whether we can identify different *types* of subjects, where types are defined by their “typical” best response behavior. Subjects who “typically” (to be defined later) best respond to a belief of type τ will be classified as a τ -type. Since we have three relevant types of beliefs (fine, coarse, uniform) this gives us three different types of subjects. To this we add a fourth type for whom we cannot say that they follow any particular type of beliefs. This type we call “other.”²³ The difference with Subsection 4.1 is that we follow each individual Row player across all 60 rounds, and we ask ourselves whether the subject’s use of feedback over time is consistent with either the fine, the coarse or the uniform belief throughout the duration of the experiment.

The use of the word “typical” in the above paragraph might sound ambiguous but in what follows we shall precisely explain our procedures. The key idea is that

²³This is justified as in almost all rounds the choices made by the first 3 types span the strategy set for row players. The “Other” type could therefore be interpreted as playing a strategy that mixes between the strategies of the other types or has another reason for stochastic choices (such as the maxmin strategy). In this sense our candidate strategies span all possible observed behaviors.

for each type of belief we calculate a variable that indicates the realized gains or losses from the actual choice compared with the respective best response to that type of belief. More specifically, when a player does not best respond to a particular belief this variable is equal to the difference between the expected profit from taking the optimal choice, minus that of the choice taken. This is the expected loss that we have also used above in the analysis of the aggregate data.

When a player does best respond, the variable is set equal to the difference between the second best choice and the best one that he did choose, which is thus negative. One might refer to this as the avoided expected loss. Thus, the variable is positive when subjects do not best respond, and negative when they do. Let's call this variable the subject's *incentive* with respect to a particular belief. I_F , I_C and I_U give the incentives with respect to fine, coarse and uniform beliefs, respectively. For each subject we, thus, have a sequence of 59 values (one for each round other than round 1) for all three incentives.

If a player typically best responds to a given belief, then we can expect two things. First, the median of the observed incentives will be negative. This is straightforward—we are just saying that the player best responds more often than not and thus that there are more negative values than positive. Second, in line with standard models of noisy decision making (see McKelvey and Palfrey, 1995) we should expect that subjects who, in principle, want to follow a particular best response mode, but are prone to mistakes, would more often deviate if the incentives were small than if they were large. (This is the same as saying that more costly errors are less likely to occur than less costly errors.) Hence, for a player who “typically” best responds against a certain belief we should expect the observed distribution of incentives *to be skewed* towards larger negative and smaller positive values.²⁴

Following this line of reasoning we define the following two-step hypothesis test:

1. Try to reject the null hypothesis that $\text{Median}(I) \geq 0$
2. If we cannot reject in part 1, then we try to reject the second hypothesis that $\text{Skew}(I) \geq 0$ (assuming that $\text{Median}(I) = 0$)

For each method of forming beliefs, we have different bounds for the incentives and different distributions under both the null hypothesis, and the different alternative hypotheses. We therefore require a test for our null hypothesis that is invariant to the scale and the particular shape of the distribution of the incentives, apart from median and skewness. We can do this by employing the following composite test. For the first part of the null hypothesis we employ a sign test, which simply relies on a binomial test statistic based on the number of positive and negative values.

²⁴We are slightly mis-using the concept of the skew of a sample in considering the relative sizes of samples on each side of the median, rather than the mean, as is more conventional. However, this interpretation is consistent with our test procedure, and all mention of the ‘skew’ of a distribution should be interpreted as using the median, rather than the mean as the location measure in its calculation.

For the second part we follow Gibbons and Chakraborti (2003) who suggests that the Wilcoxon matched-pairs rank-sum test can be used as a test for distributional symmetry (assuming a known *median*).²⁵

We proceed as follows. We take an individual subject and fix the type of belief. For belief τ this gives us, as explained above, 59 values of the incentive variable I_τ . We then use the sign test to try to reject the first part of the null hypothesis. If we are able to reject the first part of the null hypothesis, then we allocate the player to the corresponding belief type τ . If we are unable to reject it, then we employ the Wilcoxon test to see whether we can reject the second part of the null hypothesis. If we can, we again allocate the player to the belief type τ . We repeat this process for each belief type τ . If the subject cannot be allocated to any of the three types of beliefs we classify him as “other.” For all (one-sided) tests we employ on a significance level of 5%. Notice that our procedure ensures that no subject can be allocated to more than one type, unless the best responses coincide very often. This occurred in only one session of the HARD treatment.²⁶

In essence, we are allocating players to a type only if we can reject that their incentive values could be generated by a player who does not use them to make decisions. This may seem to be a test of relatively low strength, and it may seem likely that there would be high numbers of “others.” But, of course, this implies that if we can classify subjects into Fine belief, Coarse belief or Uniform belief types, we can be fairly confident in this classification. Finally, we are presented with a problem if more than one of our null hypotheses are rejected at each stage. Thankfully, this happens in only two cases out of all 64 Row players—in one particular session in the HARD treatment where the best responses to the fine and coarse feedback coincided often. For these two cases, we allocated the subjects to the type where the incentive gave the lowest p -value in our test (one was allocated to the fine-belief and one to the coarse-belief type). We have included these two players in our results, but their type allocation must necessarily be treated with less confidence than those for the other 62 Row players.²⁷

Table 4 shows our classification results.²⁸ In contrast to the fears expressed

²⁵In fact, the Wilcoxon test can also be used as a test of location (when assuming a symmetric distribution). This would open up the possibility of combining both parts of the test by using just the Wilcoxon test. However, we are reluctant to follow this path as it is not clear how the skew and location effects are compensated in the test. This would make the results ambiguous if the skew and median of a sequence of incentives had different signs. Also, we specifically want to ignore the effect of the skewness if the player best responds most of the time.

²⁶In this session, a row player chose the third row sufficiently often to influence the column player’s choices such that the Coarse ABEE equilibrium strategy was optimal for both the Fine and Coarse feedback.

²⁷We also note that the second step in our test, when reached, rejected the second part of our hypothesis in only three instances allowing us to allocate an additional three subjects to non-other types. This high level of correlation between the tests is why we didn’t apply a Bonferroni or similar correction to our multiple hypothesis test.

²⁸In COARSE, fine belief types are meaningless as row players had no access to fine statistics. Similarly, coarse belief types are meaningless in FINE. We have nevertheless kept all four types in

Type	FINE	FEASIBLE	HARD	COARSE
Fine beliefs	12	14	3	3
Coarse beliefs	1	1	7	12
Uniform beliefs	1	0	0	1
Other	2	1	6	0

Table 4: Type allocation over all rounds

above, our conservative procedure classifies the vast majority of subjects as entertaining either fine or coarse beliefs.

We find that in treatments FINE and FEASIBLE the vast majority of subjects are indeed classified as entertaining fine beliefs. Thus, what we have seen above in the aggregate is confirmed on the individual level. Subjects do use the fine information systematically and systematically best reply against it in these treatments. In COARSE the most frequent type is clearly the coarse-belief type, thereby showing that subjects mostly use the coarse feedback to form beliefs.²⁹ In treatment HARD, the evidence is less clear as there is one third of subjects classified as “other”. But, we still see a majority of subjects classified as coarse-belief type. We also see three subjects who are classified as entertaining fine beliefs. It should be mentioned here that those subjects of fine-belief type earned on average £1 more than those of coarse-belief type,³⁰ thereby suggesting that in our experiment it pays-off to have a finer understanding of the feedback.

The large number of subjects classified as “other” in HARD can be due to the fact that players need time to adjust their mode of behavior (learning explanation) or to the fact that players rely less on the information provided to them when feedback is too hard to process.³¹ To improve our understanding, we split the experiment into four phases of 15 rounds each and we repeat the same tests, again with a 5% significance level, for each phase of the experiment and each Row player, considering the possibility that each player’s method of forming beliefs (as well as the information contained in those beliefs) might change over the various phases of the experiment. This allows us to see how the number of players classified as “other” in HARD evolves through the four phases, and it also allows us to address further issues related to learning dynamics such as how subjects’ sophistication evolves with time.³² The

all treatments so as to facilitate the comparison.

²⁹In COARSE the classification as fine-belief type is meaningless as Row subjects could not materially have access to the required information.

³⁰The coarse-belief type players in turn earned just over £1 more than those of type Other, on average.

³¹In the latter case, possible modes of behavior include that players play a best-response to the uniform belief as considered above or that they themselves adopt a random strategy (either uniform strategy or maxmin strategy in a way to protect themselves against the behavior of column players).

³²Perhaps some subjects start out ignoring the feedback but are able to interpret it in some way or other later on in the session. Or perhaps some players are confused about how to interpret the information at all, and require a few rounds to understand how to use the available information to

Rounds 1 to 15				
Type	FINE	FEASIBLE	HARD	COARSE
Fine beliefs	9	9	1	1
Coarse beliefs	1	0	3	4
Uniform beliefs	0	0	0	0
Other	6	7	12	11

Rounds 16 to 30				
Type	FINE	FEASIBLE	HARD	COARSE
Fine beliefs	12	12	3	2
Coarse beliefs	0	0	5	7
Uniform beliefs	1	0	0	1
Other	3	4	8	6

Rounds 31 to 45				
Type	FINE	FEASIBLE	HARD	COARSE
Fine beliefs	12	13	3	2
Coarse beliefs	1	0	8	10
Uniform beliefs	0	0	0	1
Other	3	3	5	3

Rounds 46 to 60				
Type	FINE	FEASIBLE	HARD	COARSE
Fine beliefs	12	12	3	1
Coarse beliefs	2	0	9	13
Uniform beliefs	0	0	0	1
Other	2	4	4	1

Table 5: Type allocation over different phases

results are shown in Table 5.

Initially, we find a large share of Other types but after the first quarter of the experiment there is a huge jump in sophistication. From the second quarter onwards most subjects entertain either fine or coarse beliefs and then keep on doing so for the rest of the experiment. These results suggest that subjects in the HARD treatment are learning to process the information during the experiment, and sometimes change how they process the feedback during a session.³³ Interestingly, it also seems that, once subjects have settled on a way of accessing and interpreting the given information, they rarely change this interpretation later in the experiment.³⁴ Thus,

make decisions. Or it could happen that subjects after a while find it too tedious to pay attention to the information provided and use simpler choice rules.

³³This provides evidence that the subjects are not ignoring the feedback to play a maxmin or uniform strategy.

³⁴Among the three fine belief types in the third quarter of HARD, two remained so in the last quarter (the third one switched to type other). At the same time, only one coarse belief type switched to a fine belief type while all others remained coarse belief types. In all, 25 subjects kept

in our experiment, there is no evidence that subjects disregard the information when it is hard to process, and there is no evidence either that subjects switch to more sophisticated rules of forming beliefs with more experience. The latter insights suggests that learning the strategy of the opponent is a separate process from learning how to interpret feedback or, more generally, learning how to learn to play the game.

6 Related literature

Our paper belongs to the family of experimental papers concerned with the understanding of learning processes. The papers from that strand most closely related to ours are by Cooper and Kagel (2007ab) who focus on cross-game learning, the ability of subjects to take what has been learned in one game and transfer it to related games. 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. The common feature with our own experiment is that several games are being considered. Yet, there are several important differences between our investigation and theirs. First, our main message is about the need to extend the notion of equilibrium beyond Nash equilibrium in the presence of learning spillovers whereas Cooper and Kagel focus on whether cross-game learning may facilitate learning to play a Nash equilibrium in a new but related game. Second (and in close relation to the first difference), the various games are considered simultaneously in our setup and sequentially in theirs. It may also be mentioned that in our setup the games are not related in any clear way (the best-response correspondences are very different in the two games) and thus the finding of learning spillovers in our setup should be considered as more surprising than in setups in which games are more closely related.

There are obviously several other strands in the experimental literature on learning that we now discuss in relation to our experiment. First, Erev and Roth (1998) and Camerer and Ho (1999) examine the important issue which of the subject's own past performance or opponents' past play explain better the learning dynamics. While the early contributions in this strand have suggested that reinforcement learning explains better the data, this has been recently challenged in Wilcox (2006). In our experiments 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.³⁵

the same type for all phases and a further 27 were type-Other before changing to, and remaining as, one of the other types.

³⁵It would, of course, be nice to see what happens if subjects get to know their past performance in addition to the feedback they currently receive. But, remember that our primary interest lies in the study of learning spillover, hence explaining our choice of experimental design as a first step.

We should also mention that most experiments concerned with how subjects form beliefs about their opponent's behavior also inform subjects about their opponent's payoff structure, which opens the door to further inferences based on introspective reasoning about the opponent's incentives (see Ehrblatt et al. (2005) for an explicit account of this).³⁶ By not providing the opponent's payoff structure, we somehow force the belief to be solely driven by opponents' past behaviors, which again allows us to better focus on learning spillover.³⁷ In the same strand of literature, some papers have developed the methodology of eliciting subjects' beliefs (about opponents' next play) at the same time as learning takes place. While Nyarko and Schotter (2002) find that subjects tend to play best-response to their announced beliefs, Costa-Gomes and Weizsäcker (2006) challenge this conclusion in their experiment. All these studies are silent though on how beliefs are formed in the first place. By contrast, in our experiment, we did not elicit beliefs but we postulated that beliefs are either the distribution of opponents' actions in the game to be played (fine beliefs) or the aggregate distribution over the two games (coarse beliefs), and we suggested that one or the other method of forming beliefs provides a good account of observed behaviors depending on the treatment.

Our investigation has also some connection with the experimental literature interested in the cognitive abilities of subjects (see Stahl 1993, Nagel 1995, Costa-Gomes et al. 2001 and Camerer et al. 2004), since in treatment HARD, we observe some (few) subjects who are better able of processing information than most others toward the end of the experiment (see section 5). Yet, an essential difference is that this literature is mostly concerned with non-repeated play so that there is no scope for learning. Besides, from the viewpoint of this literature, only levels 0 and 1 are meaningful in the present context, since subjects were not informed of their opponent's payoff structure and had thus very limited opportunity to reason about their opponent's incentives.

Finally, our experiment has also some connection with those experiments that try to understand which games players perceive to be playing.³⁸ 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 learning spillover as experimentally explored in our paper is not merely of viewing our subjects as playing a Nash equilibrium of a different game but rather as playing a different equilibrium of the original (set of) games.

³⁶Oechssler and Schipper (2003) is a notable exception to this literature as they study 2x2 games where subjects are not informed about their opponents' payoffs. Their focus is, however, on the question whether through repeated play subjects can infer their opponents' best-reply correspondence.

³⁷From the viewpoint of fictitious play learning models such as those described in section 2, it is irrelevant whether or not players get to know their opponent's payoff structure. We suspect though that such models are more likely to be relevant if this extra information is not available.

³⁸See Devetag and Warglien (2007) for a recent experiment on this (and also Oechssler and Schipper (2003) for another perspective).

7 Conclusion

Real life, one might argue, is a collection of multiple games, some repeated, others not. Many of these games will be quite similar in structure. For example, take the games we play with different sets of colleagues, which may involve subtle yet crucial differences. Specifically, one may or may not be cooperative when a colleague asks for a favour one is not contractually obliged to do. But the response in that kind of situation may (in principle) crucially depend on whether this colleague works in the same or a different division or whether he is of the same, a lower or a higher rank. However keeping track of the responses in these different games may be a difficult task and it seems plausible that many people rather keep track of the aggregate response *across* a number of different games.

From a practical viewpoint, it seems clear that to process information and make decisions, we use simplifying heuristics all the time.³⁹ Think, for example how one judges academic CVs that, after all, provide very complex information. A first approximation might simply be to count publications in good journals, neglecting time dimension or even titles. Such simplifying heuristics may in turn give rise to learning spillovers of the kind we obtain here.

More generally, whether or not mental accounting is precise and game-specific or coarse and aggregated across games may be a function of the structure and framing of information. Intuitively, the more accessible and precise information is, the easier it will be to keep track of the behaviors in the different games separately. The less accessible, the more likely are (learning) spillovers from one game to the other.

In this paper we have taken two abstract games to analyze whether and how such spillovers occur. Starting with two extreme benchmarks of exogenous easy-to-access precise or coarse information we have ventured into rather unknown territory, gradually changing the accessibility of precise information. In doing so we have chosen frames that lend themselves quite naturally to simple heuristics that bundle information. Our main results were that learning spillovers do occur and that their occurrence does indeed depend on the accessibility of feedback information. Moreover, it turns out that these spillovers *can be modelled*—and in quite familiar ways using the concept of analogy-based expectation equilibrium. Thus, acknowledging the role framing of information and learning spillovers does not imply giving up the apparatus of game theoretic modeling and equilibrium analysis. We consider this as good news.

Future work should focus on the kind of learning heuristics used by agents in real economic interactions so as to better predict whether and when learning spillover should be expected and in which form. We note that beyond the objective of better describing the real world, the understanding of accessibility and learning spillovers may also be of primary importance from a normative viewpoint (i.e. from a mechanism design perspective) as a mechanism that functions only if agents extract

³⁹Somehow the work of Aragones et al. (2005) provides some rationale for this.

information they cannot be realistically expected to extract may fail horribly in practice.⁴⁰

References

- [1] Aragoes, E., I. Gilboa, A. Postlewaite and D. Schmeidler (2005): “Fact-free Learning,” *American Economic Review* **95(5)**, 1355–1368.
- [2] Camerer, C., and T.H. Ho (1999): “Experience-weighted Attraction Learning in Normal Form Games,” *Econometrica* **67**, 827–874.
- [3] Camerer, C., T.H. Ho and J.K. Chong (2004): “A Cognitive Hierarchy Model of Games,” *Quarterly Journal of Economics* **119(3)**, 861–898.
- [4] Cooper, D.J. and J.H. Kagel (2007a): “Learning and Transfer in Signaling Games,” mimeo.
- [5] Cooper, D.J. and J.H. Kagel (2007b): ”The Role of Context and Team Play in Cross-Game Learning,” mimeo.
- [6] Costa-Gomes, M., V. Crawford, and B. Broseta (2001): ”Cognition and Behavior in Normal-Form Games: An Experimental Study,” *Econometrica* **69**, 1193-1235
- [7] Costa-Gomes, M. and G. Weizsäcker (2006): “Stated Beliefs and Play in Normal-Form Games,” mimeo LSE.
- [8] Devetag, M.G. and M. Warglien (2007): “Playing the Wrong Game: An Experimental Analysis of Relational Complexity and Strategic Misrepresentation,” forthcoming *Games and Economic Behavior*.
- [9] Ehrblatt, W.Z., K. Hyndman, E.Y. Ozbay and A. Schotter (2005): “Convergence: An Experimental Study,” mimeo NYU.
- [10] Ellsberg, D. (1961): “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics* **75**, 643–669.
- [11] Erev, I., and A.E. Roth (1998): “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *American Economic Review* **88(4)**, 848–881.

⁴⁰Of course, such behavioral mechanism design needs a profound understanding of the mapping from frames into outcomes for which further research is required. Some interesting results on the role of strategies’ labels and context can be found in Cooper and Kagel’s (2007ab) study on cross-game learning. They find that subjects are better able to transfer what they have learned in one situation to a new situation if meaningful real-life (here: market) context is provided. A related finding is reported in Huck, Normann, and Oechssler (2004) where a market frame helps subjects to understand and overcome the dilemma problem imposed in 2-player Cournot games.

- [12] Ettinger, D. and P. Jehiel (2006): “Towards a Theory of Deception,” mimeo PSE, U. Cergy-Pontoise and UCL.
- [13] Gibbons, J.D., and S. Chakraborti (2003): *Nonparametric Statistical Inference*, Marcel Dekker.
- [14] Higgins, E. T. (1996): “Knowledge Activation: Accessibility, Applicability, and Salience,” in E. Tory Higgins and Arie W. Kruglanski, eds, *Social Psychology: Handbook of basic principles*. New York: Guilford Press, pp 133–168.
- [15] Huck, S., H.-T. Normann and J. Oechssler (2004): “Two are few and four are many,” *Journal of Economic Behavior and Organization* **53**, 435-446.
- [16] Jehiel, P. (2005): “Analogy-based Expectation Equilibrium,” *Journal of Economic Theory* **123**, 81–104.
- [17] Jehiel, P. and F. Koessler (2007): “Revisiting Games of Incomplete Information with Analogy-based Expectations,” forthcoming *Games and Economic Behavior*.
- [18] Khaneman, D. (2003): “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review* **93**, 1449–1475.
- [19] Knight, F.H. (1921): *Risk, Uncertainty, and Profit*, Boston and New York: Houghton Mifflin.
- [20] McKelvey, R.D. and T.R. Palfrey (1995): “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior* **10**, 6–38.
- [21] Milgrom, P., and J. Roberts (1982): “Predation, Reputation, and Entry Deterrence,” *Journal of Economic Theory* **27**, 280–312.
- [22] Milgrom, P., and J. Roberts (1991): “Adaptive and Sophisticated Learning in Normal Form Games,” *Games and Economic Behavior* **3(1)**, 82-100.
- [23] Nachbar, J. (1990): “‘Evolutionary’ Selection Dynamics in Games: Convergence and Limit Properties,” *International Journal of Game Theory* **19**, 59–89.
- [24] Nagel, R. (1995): “Unraveling in Guessing Games: An Experimental Study,” *American Economic Review* **85**, 1313–1326.
- [25] Nyarko, Y., and A. Schotter (2002): “An Experimental Study of Belief Learning Using Elicited Beliefs,” *Econometrica* **70(3)**, 971–1005.
- [26] Oechssler, J., and B. Schipper (2003): “Can you guess the game you’re playing?,” *Games and Economic Behavior* **43**, 137-152.

- [27] Stahl, D.O. (1993): “Evolution of smartn Players,” *Games and Economic Behavior* **5**, 604–17.
- [28] Wilcox, N.T. (2006): “Theories of Learning in Games and Heterogeneity Bias,” *Econometrica* **74(5)**, 1271–1292.

Appendix

A Experiment Instructions

There follows the instruction sheets handed out at the start of the experiment. Step-by-step instructions were presented on the computer screen at the start of each session, and these sheets were intended to be used as a reference during the experiment. They have all been amended to show the name of the treatment they apply to, but are otherwise identical to those handed out.

ELSE Experiment
Instructions summary sheet
(Column Participant – All Treatments)

You are a COLUMN participant. Choose your COLUMN using the buttons on the matrix below.

Situation A

	C1	C2	C3	C4	C5
4	Payoff 1	Payoff 2	Payoff 3	Payoff 4	Payoff 5
2	Payoff 6	Payoff 7	Payoff 8	Payoff 9	Payoff 10
4	Payoff 11	Payoff 12	Payoff 13	Payoff 14	Payoff 15

This is how many times the row participants have chosen row 1 in situation A in the last 5 rounds

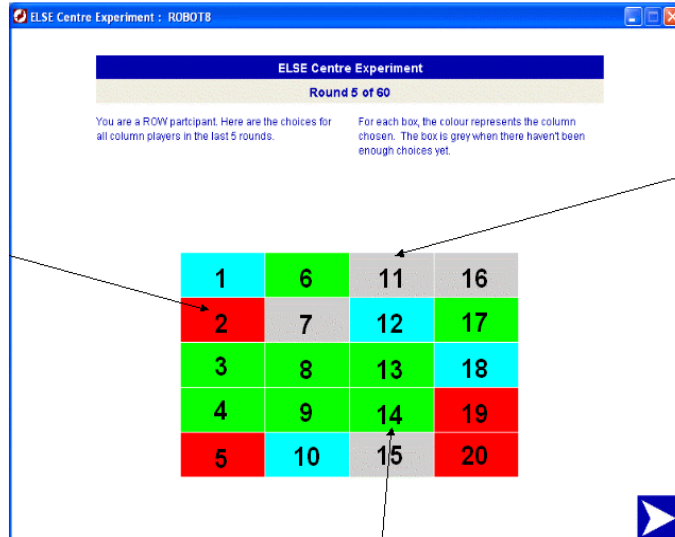
Press this button to choose column 3

This is how much you get paid for this round if you pick column 4 and the row participant you're matched with picks row 3.

ELSE Experiment

Instructions summary sheet

(Row Participant – Hard Treatment)



A column participant chose the red column in the last 5 rounds. The situation was the 2nd in the sequence on the next screen.

Grey boxes mean not enough choices have been made yet.

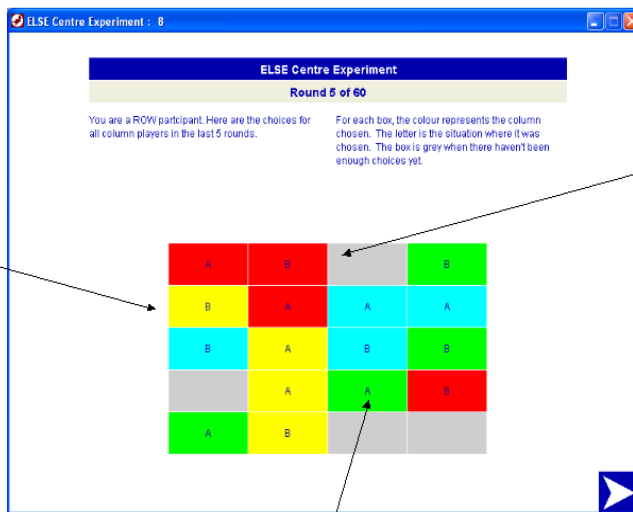
A column participant chose the green column in the last 5 rounds. The situation was the 14th in the sequence on the next screen.

6

ELSE Experiment

Instructions summary sheet

(Row Participant – Feasible Treatment)



A column participant chose the yellow column in situation B in the last 5 rounds.

Grey boxes mean not enough choices have been made.

A column participant chose the green column in situation A in the last 5 rounds.

5

ELSE Experiment

Instructions summary sheet

(Row Participant – Feasible and Hard Treatments)

Press this button to choose row 1

R1	25p	0p	10p	0p	0p
R2	20p	15p	5p	10p	0p
R3	15p	10p	0p	5p	25p

Each choice for the column participant has a different colour

This is how much you get paid for this round if you pick row 3 and the column participant you're matched with picks column 4.

4

ELSE Experiment

Instructions summary sheet

(Row Participant – Fine Treatment)

Press this button to choose row 1

	3	1	1	3	2
R1	Payoff 1	Payoff 2	Payoff 3	Payoff 4	Payoff 5
R2	Payoff 6	Payoff 7	Payoff 8	Payoff 9	Payoff 10
R3	Payoff 11	Payoff 12	Payoff 13	Payoff 14	Payoff 15

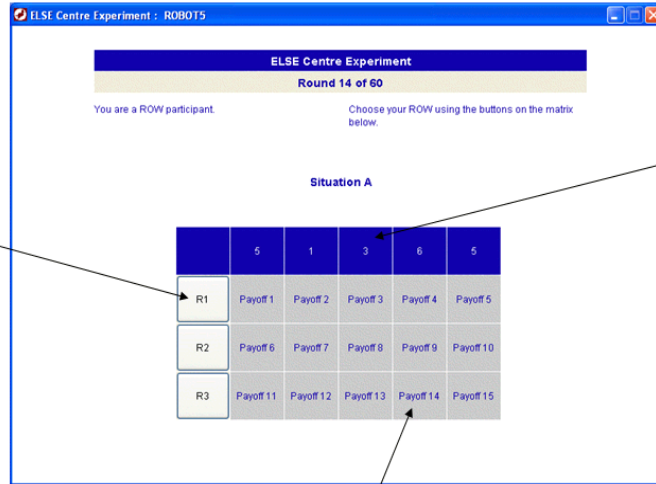
This is how many times the column participants have chosen column 3 in situation A in the last 5 rounds

This is how much you get paid for this round if you pick row 3 and the column participant you're matched with picks column 4.

ELSE Experiment

Instructions summary sheet

(Row Participant – Coarse Treatment)



Press this button to choose row 1

This is how many times the column participants have chosen column 3 in either situation in the last 5 rounds

This is how much you get paid for this round if you pick row 3 and the column participant you're matched with picks column 4.