

Equilibrium Selection through Incomplete Information in Coordination Games: An Experimental Study

Antonio Cabrales, Rosemarie Nagel and Roc Armenter^{*}

February 15, 2002

^{*} Contact: Antonio Cabrales, Universitat Pompeu Fabra, antonio.cabrales@econ.upf.es; Rosemarie Nagel, Universitat Pompeu Fabra, rosemarie.nagel@econ.upf.es; Roc Armenter, Northwestern University, armenter@nwu.edu. We thank Urs Fischbacher for teaching us to program this experiment and for providing his program, z-tree, Martin Menner for writing the code for the simulations, Frank Heinemann, Nick Vriend and seminar audiences at UPF, Munich, Stockholm for helpful comments. Cabrales and Nagel acknowledge the financial support of Spain's Ministry of Education under grant BEC2000-1029 and PB98-1076 and the Generalitat de Catalunya under grant 1999SGR-00157.

Abstract: We perform an experiment on a pure coordination game with uncertainty about the payoffs. Our game is closely related to models that have been used in many macroeconomic and financial applications to solve problems of equilibrium indeterminacy. In our experiment each subject receives a noisy signal about the true payoffs. This game has a unique strategy profile that survives the iterative deletion of strictly dominated strategies (thus a unique Nash equilibrium). The equilibrium outcome coincides, on average, with the risk-dominant equilibrium outcome of the underlying coordination game. The behavior of the subjects converges to the theoretical prediction after enough experience has been gained. The data (and the comments) suggest that subjects do not apply through “a priori” reasoning the iterated deletion of dominated strategies. Instead, they adapt to the responses of other players. Thus, the length of the learning phase clearly varies for the different signals. We also test behavior in a game without uncertainty as a benchmark case. The game with uncertainty is inspired by the “global” games of Carlsson and Van Damme (1993).

Keywords: Global games, risk dominance, equilibrium selection, common knowledge.

JEL Classification: C72, C91, C92, D82, G10, G21.

1. INTRODUCTION

There are many important economic phenomena for which a natural model is a coordination game, that is, a game with multiple equilibria. Often, these equilibria are Pareto-ranked. For example, unemployment or underemployment (Diamond 1982, Hart 1982, Bryant 1983), currency crises (Obstfeld 1996), bank runs (Diamond and Dybvig 1983)¹ have all been modeled this way. Some of the equilibria in these models look more plausible, and the policy implications also depend on which equilibrium is the one that more closely represents reality.

But, without further information, it is hard to know what determines which equilibrium will be observed. In principle, even a sunspot, rather than any fundamental information, can be the explanation. This is not the way that people involved in these games usually perceive the situation² and theoretically it is not very satisfactory. Furthermore, with multiplicity of equilibria, even comparative statics exercises are not straightforward (a change of policy parameters could, for example, trigger a change of equilibrium).

A recent approach to this problem has been to introduce incomplete information in the model. In this way a unique equilibrium can arise. Often, this equilibrium is the one that was thought to be the more economically plausible one. What is perhaps surprising is that *even small departures from complete information* (“almost common knowledge”) can generate unique equilibria, when under common knowledge there is more than one. This is already clear in the e-mail game of Rubinstein (1989)³. Carlsson and Van Damme (1993) “slightly” perturb the players’ information in a coordination game.⁴ This leads to an incomplete information game, which has a unique solution by iterated deletion of *strictly* dominated strategies. Morris, Rob and Shin (1995), and Kajii and Morris (1997) further explain and generalize the logic behind the result of Carlsson and Van Damme (1993).

Many researchers have applied this technique to solve the equilibrium indeterminacy in the important economic situations we mentioned above. For example, Burdzy, Frankel and Pauzner (2001), and Frankel and Pauzner (2001) apply this technique in dynamic settings. Goldstein and Pauzner (1999) do it with bank runs and Goldstein (1999) for banking and currency crises. Heinemann (1999), Heinemann and Illing (1999), and Morris and Shin (1998),

¹ Cooper and John (1988) show that many of these models share the feature that there are spillovers between different players’ strategies. These spillovers create strategic complementarities, which are the root for the multiplicity of equilibria.

² Most newspaper accounts of the recent currency crisis (with an added bank run element) in Argentina focus on fundamental factors, rather than on sunspots.

³ Other related examples can be found in Myerson (1991, p.66) or Binmore (1992, p.445).

⁴ These games have two Nash equilibria in pure strategies, and one mixed-strategy equilibrium.

deal with speculative attacks. Morris, Postlewaite and Shin (1995) study bubbles; and Shin (1996) asset trading..

Despite their theoretical and practical attractiveness, these ideas face some potential empirical difficulties. The logic of the result involves that players do many rounds of deletion of dominated strategies, each of which requires sophisticated Bayesian reasoning. It is well known that real-life players are not particularly good at performing either of these tasks. Sefton and Yavaş (1996), for example, study experimentally a game that is derived from a coordination game. The unique solution of the transformed game, which also involves many rounds of deletion of strictly dominated strategies, is not the most observed experimental outcome⁵. More generally, experiments have shown that subjects apply low levels of reasoning (Stahl and Wilson 1994, Nagel 1995, Costa-Gomes, Crawford, and Broseta 2001, Costa-Gomes, Crawford 2002). In this paper, we conduct an experimental test of a game inspired in Carlsson and van Damme (1993).

We find that, if the game is played repeatedly, the theoretical solution is the most observed outcome empirically, when enough rounds elapse. Furthermore, the frequency of play for actions whose optimality does not require assumptions about other agents' rationality decreases fast. Actions that are optimal only if other (dominated) actions are played infrequently, decrease their frequency of play at a slower speed. However, our subjects clearly do not arrive at the solution by careful introspection⁶, they rather adapt their play in view of their observations about the other players' behavior.

Unlike in Sefton and Yavaş (1996), the theoretically predicted outcome in our game is the risk dominant solution (Harsanyi and Selten 1988) in the underlying complete information game. This confirms the empirically usefulness of risk dominance as an equilibrium selection criterion, which is documented with experiments on coordination games with complete information as e.g. in Van Huyck, Battalio and Beil (1990, 1991), Friedman (1996), Charness (1998), Cabrales, García-Fontes and Motta (2000), Charness and Grosskopf (2000) and also in our benchmark treatment with complete information.

Kajii and Morris (1997) actually show that a sufficient condition for robustness to incomplete information is a many-player many-action generalization of risk dominance, which suggests that our results may generalize well beyond the setup in this paper.

⁵ In a betting game by Sonsino, Erev and Gilat (1999) Sovik (2001), which is also solved by iterated deletion of strictly dominated strategies, the empirical solution does not coincide with the theoretical solution. It is possible, though, that the number of periods was not enough for equilibrium convergence, as the trend in her data suggests.

⁶ Both their actions and their responses in the comments' sheets suggest that they do not grasp the subtle Bayesian issues involved.

The remainder of the paper is structured as follows: in section 2 we discuss theoretically the game that is played by the experimental subjects. Section 3 describes the experimental design. Section 4 collects the experimental results. These results are discussed in Section 5. Section 6 concludes.

2. THE GAME

The experimental subjects face an incomplete information game in which there are 2 players, each of whom has 2¹ actions. The payoffs for each action profile are given in the following matrix

	<i>A</i>	<i>B</i>
<i>A</i>	X,X	X,0
<i>B</i>	0,X	80,80

True value X	90		80			70			60			50	
Signal S	T	U	T	U	W	U	W	Y	W	Y	Z	Y	Z
P(X)	1/5		1/5			1/5			1/5			1/5	
P(S X)	1/2	1/2	1/3	1/3	1/3	1/3	1/3	1/3	1/3	1/3	1/3	1/2	1/2
P(X S)	3/5	3/7	2/5	2/7	1/3	2/7	1/3	2/7	1/3	2/7	2/5	3/7	3/5
P(S _i S _j)	P(T W)=P(U W)=1/3; P(U Y)=P(W Y)=2/7; P(W Z)=P(Y Z)=1/3												

Signal S	T	U	W	X	Z
U(A S)	86	81.429	70	58.571	54
U(B S)*	<80	<80	<53.333	<57.142	<34.666
Iterated reasoning	Step 1 (elimination of strictly dominated strategy B)		Step 2	Step 3	Step 4

Table 1a and b: True Values, Signals, conditional probabilities and expected utilities

* For the expected utility U(B|S) at step i, we assume that all individuals on step j<i never play B and those on step j>i play at best always B.

The random variable X can take 5 possible values, 90, 80, 70, 60 and 50 (row 1 in table 1a). Each of these variables have the same probability $p=1/5$ (row 3 of table 1a). The players do not know the value of X, but independently receive a private signal about the value of X. The signals S_i can be: T, U, W, Y, Z. The relationship between signals S and the value of X are shown in row 1 and 2 of table 1a together with the conditional probabilities P(S|X) in row 4. Furthermore, with this structure, we can construct conditional probabilities P(X|S) (shown in row 5). The resulting expected utilities for each action C={A,B} conditional on the signal, U(C|S), are shown in table 1b. We assume common knowledge of rationality (see also proof in the appendix). In the last row of table 1b we indicate the number of rounds of elimination needed, given a signal S, so that playing action is strictly dominated.

The strategy for all players is a function from the set of signals to the set of actions.

Proposition: The only strategy in this game that survives the iterated deletion of strictly dominated strategies is to play action A irrespective of the signal received.

Proof: for complete proof, see appendix.

Let us give some intuition here. The proof is done in four steps of iterated reasoning.

Step 1 (elimination of strictly dominated strategy): A rational player with signal T or U will never play B, since the expected utility of action A, is greater than for action B, irrespectively what a player with signal W or Y does.

Step 2: Knowing that a player with signal T or U will play A, a rational player with signal W will never play B, irrespectively what a player with signal W, Y or Z does.

Step 3: Knowing that players with signal T, U, or W will never play B, a rational player with signal Y will never play B, irrespectively what a player with signal Y or Z does.

Step 4: Knowing that players with signal T, U, W, or Y will never play B, a rational player with signal Z will never play B, irrespectively what a player with signal Z does.

The signal received by the agents in Carlsson and Van Damme (1993) is distributed continuously in a small interval around the value of X . We chose to modify the game in order to make the theoretic reasoning simpler. In particular, one can see in the proof of the proposition that the equilibrium can be reached after only four rounds of deletion of dominated strategies. Even so, the task still seems to be challenging.

It is also worth noting that the proposition assumes that the figures in the payoff matrix are Von Neumann-Morgenstern utilities, however we do not control for risk preferences in the laboratory. It is not difficult to see that, if agents are risk loving, one can find other equilibria in this game in which an action is signal dependent. We could have tried to control for risk preferences, but Selten, Sadrieh and Abbink (1995) have cast some doubts on procedures to control for risk aversion. Our results can be better explained by assuming that some agents are mildly risk loving and do not have a good assessment about play probabilities at the beginning of the game.

3. EXPERIMENTAL DESIGN

The experiments were run with undergraduates of all faculties in the Leex of the Universitat Pompeu Fabra. No subject could participate in more than one session. Upon arrival students were randomly assigned to their seats. One of the instructors read the instructions (see appendix) aloud and questions were answered in private.

We run three different treatments with different time horizons (15 periods and 50 periods) and/or certainty or uncertainty of the payoffs. The treatment with uncertainty and 50 periods was added because behavior did not converge after 15 periods.

- UC15: In the first treatment each player faced the payoff matrix shown in section 2 with uncertainty about the value. Subjects played 15 rounds of the same game. We will call this treatment "*uncertainty15*" or "*UC15*". The value X could be 50, 60, 70, 80 or 90, where the realization was unknown to the subjects when making a decision. However, each subject received a signal about X as private information. The relationship between the value X and the signal as shown to the subjects is given in the first two rows in table 1. In the instructions we mentioned the probabilities of 20% for each value X . Furthermore, we stated that each possible signal for a given X had equal probability without specifying these probabilities. We did not state probabilities of X conditional on these signals. We informed the players that the signal for each player was from the same distribution and independently drawn. We emphasized various times that the underlying variable X was the same within a pair but the signal can be different.
- UC50: The second treatment was identical to the first except the number of rounds was 50. We will call this treatment "*uncertainty50*" or "*UC50*".
- C15: The third treatment is without uncertainty in the payoff matrix. Instead the value of X is always restricted to 50, that is the lowest value of all possible X s in the uncertainty treatments. The number of rounds is 15. We call this treatment "*certainty15*" or "*C15*".

The third treatment serves as a benchmark treatment, in order to check which choice is predominant in the complete information case. The lowest value of X was chosen, since the signals of that value most likely induced a B choice in the uncertainty treatment.

After each period each player was informed about the choice of the other player and both players' payoffs. In the uncertainty treatments he received the information about the true value of X . The signal of the other player was never revealed.

The first and third treatments were run by hand, and the second by computer using z-tree-tools. In all sessions there were 16 subjects. Each subject was randomly matched with an opponent in each round. However, in order to obtain two independent observations in each session, we only matched the first 8 subjects with each other and the last 8 subjects with each other. In the hand experiments the matching was the same for each session and also the randomly chosen signals for the two independent observations were maintained in each session. In the first treatment we run 3 sessions (6 independent observations) and in the second and third treatment we have 2 sessions with 4 independent observations each. All sessions lasted about 1 hour including reading the instructions.

4. EXPERIMENTAL RESULTS

The main points of this study are to test whether the introduction of payoff uncertainty in a coordination game leads to the unique equilibrium of the incomplete information game and whether different signals induce a different speed of convergence since different steps of reasoning are required. A further question is whether this convergence is faster than in studies with complete information as in our certainty15 treatment.

In figures 1a,b (see appendix) we show the relative frequencies of B choices over all rounds conditional on each possible signal, separately for each UC-treatment (figure 1a with 15 rounds for uncertainty15, and figure 1b, 50 rounds for uncertainty50). Clearly, the higher the step of reasoning, the higher the relative frequency to choose B. With signal Z this amounts to 52% and 33%, for UC15 and UC50, respectively. For T or U this amounts to 1.7% and 3.5% for the UC15 and UC50, respectively. The binomial test of equal probability of increase or no-increase of the relative frequencies with increasing depth of reasoning (from signal T to signal Z) can be rejected at a 5.5% level, with 6 independent observations for treatment UC15 and 4 for UC50. The relative frequencies for B-play (dominated strategies) given T or U are similar as in other studies on iterated reasoning (see e.g. Sovik (2000)).

The dynamics of behavior over time is shown in figures 2a and b. Here we separate the relative frequency of B-choices in each period for each possible signal, pooled over all sessions per treatment. Subjects learn quickly to avoid a B choice given a signal T, U (after about 15 periods), or W (after 30 periods). This is a little slower with a Y signal (after 40 periods). However, a Z signal triggers some players to choose B until the 50 period, although these are a few subjects. Do subjects apply the concept of iterated elimination of dominated strategies? From the comments we do not get any indication that they do. Rather it seems that they learn to avoid B-choices when they have received zero payoffs.

In order to compare the convergence behavior under uncertainty with that under a treatment of certainty we also run a treatment with certainty of all payoffs and $X=50$. We chose $X=50$ since this is the lowest possible value X of our uncertainty treatment and provoked the most B-play. The result is that in the certainty treatment the average B-play is 49% while for UC15 and UC50, given a Z-signal it is 56% (just taking the first 15 rounds of all UC-treatments). Looking at independent observations, there is no difference in behavior (see also figure 3 where we plot the % of B-play if signal is Y or Z (aggregated over all UC-observations) and the % of B-play in the certainty treatment (C) in each of the 15 periods).

To sum up: players have to learn to play the A-choice and thus the equilibrium as proposed by Carlsson and Van Damme. For Signals with "high" expected value of choice A (or low level of reasoning) subjects learn fast to avoid the B-choice while for the signal with the lowest expected value of choice A (or high level of reasoning) some need as many as 30 periods to always choose A. Clearly, players do not seem to apply "a priori" reasoning and common knowledge of rationality to achieve the iterated elimination of strictly dominated choices but rather learn through payoff feedback to avoid B-choices.⁷ So in the next section we explore more thoroughly whether learning is a good explanation for the observed behavior and what the driving forces are for such a learning process.

5. LEARNING ANALYSIS

In this section we want to analyze the driving forces in the learning process, the fast convergence in states of high signal (T and U) and the much lower convergence for states with low signals (Y and Z). Most learning models, under very mild assumptions about behavior, predict that agents' actions converge to the unique equilibrium in this game. Samuelson and Zhang (1992) discuss the "dynamic elimination" of strategies that are strictly dominated (iteratively) for continuous time, and Cabrales and Sobel (1992) do the same for discrete time⁸. It is also more or less obvious from a casual inspection of the data (see figure 2a and b) that our subjects were quickly learning to avoid dominated strategies (playing A, given a signal T or U). As time went by, the frequency of strategies that are dominated only after some others have been eliminated (playing A, given signal W, Y or Z) also decreased.

Behavior in the long run can therefore be well explained by many learning models. However, short run behavior in this game cannot be reproduced easily by some of the simpler learning models. Therefore one needs a model with several parameters to obtain an adequate fit. We will first describe a general, flexible model, which encompasses some of the simpler ones. We will see that with some "intermediate" parameterization, the model captures well the stylized features of the short-run behavior of the data, while with simple models, this is not possible.

⁷ Carlsson and Van Damme (1993) already note this, as they say: "This [common knowledge of rationality] justification of our solution concept, however, may not be totally convincing since a major motivation for our approach has been the wish to relax the common knowledge assumption. Hence it would be desirable to find alternative justifications which [...] dispense with it altogether. The kind of stories that naturally come to one's mind are those where the strategy choices, instead of being determined by strictly rational considerations, result from some learning or evolutionary process.

⁸ Some other papers have studied this issue, eg.: Nachbar (1990) and Milgrom and Roberts (1991).

We simulate behavior with the model of Camerer and Ho (1999), which encompasses many other important learning models⁹. Let us introduce some notation to present it.

We have a 2-player game, where for each signal S , each player, i , has 2 possible actions. Let a^j_i be the action j of player i , and $a_i(t)$ denote the action played by player i at time t . The vector $a(t) = \{a_1(t), a_2(t)\}$, is the action profile of both players at time t , and $a_{-i}(t)$ is the action of the player other than i at time t . The payoff function, for a player i using action j , against an action profile a_{-i} of the other players, given a true value X , is denoted by $\pi_i(a_i^j, a_{-i} | X)$

In this model the probability that for signal S_i , the action a^j_i is chosen by player i at time $t+1$, conditional on $a(t), \dots, a(1), a(0)$ is given by

$$P_i^j(t+1 | S_i) = \frac{(A_i^j(t | S_i))^{\lambda_i}}{\sum_{k=1}^{m_i} (A_i^k(t | S_i))^{\lambda_i}}$$

In this expression, $A_i^j(t | S_i)$ is the ‘‘attraction’’ of action j for agent i for signal S_i at time t . Let the true value of X at time t be $X(t)$. Then $A_i^j(t | S_i)$ is given by:

$$A_i^j(t+1 | S_i) = \frac{\phi N_i(t) A_i^j(t | S_i) + [\delta + (1 - \delta) I(a_i^j, a_i(t))] \pi_i(a_i^j, a_{-i}(t) | X(t))}{N(t+1)}, \text{ if agent } i$$

received signal S_i at time t .

$$A_i^j(t+1 | S_i) = A_i^j(t | S_i) \text{ if agent } i \text{ did not receive signal } S_i \text{ at time } t.$$

In this formula $I(x,y)$ denotes the indicator function which is 0 if $x \neq y$ and is 1 if $x = y$, and $N(t+1)$ is a variable that is used to express the importance of past experience and is recursively defined by $N(t+1) = \rho N(t) + 1$. The variables $N(t)$ and $A_i^j(t)$ are started with some initial values $N(0)$ and $A_i^j(0)$. We thus have 10 attractions to possibly update in our simulations: Attractions of action A given signal T, U, W, Y or Z, or attractions of action B, given signal T, U, W, Y or Z.

⁹ This includes models like fictitious play (Brown 1951, Robinson 1951), best response dynamics (Cournot 1971, Matsui 1991) or the learning by reinforcement model (Bush and Mosteller 1951, Cross 1973, Roth

In order to understand the meaning of the different elements of the model and the coefficients, notice first that the attractions are related to the payoffs. This means that when λ is large, small differences in (expected) payoffs lead to large differences in the probabilities that strategies are played. For this reason, large values of λ will make the model be more like fictitious play or best response, which put weight only on the “best” strategy, rather than learning by reinforcement, which puts weight roughly in proportion to payoffs. Strictly speaking, the model is exactly like fictitious play only if $\lambda = \infty$, but there will not be much difference in behavior for high values of λ .

The parameter δ regulates how much more the payoff of the strategy that has actually been played in a given period gets incorporated in the attraction with respect to strategies that have not been played. With $\delta = 0$ only the strategy that has actually been played gets its payoff incorporated in the attractions (as it happens, for example, in learning by reinforcement). With $\delta = 1$ (as in best reply or fictitious play) the payoffs to all strategies (given the strategy played by the opponent in the present period) get incorporated in the attractions.

The parameters ρ and ϕ tell us how much the past is discounted when updating attractions. When $\rho = 1$ and $\phi = 1$, we have a model like fictitious play which gives all periods the same weight, whereas $\rho = 0$ and $\phi = 0$, is more like best reply which reacts only to last period's experience.

Our benchmark simulation for the model assumes $N=1$, $\rho = 0$, $\phi = 0.9$, $\delta = 0.5$, and $\lambda = 1.5$. As initial attractions we have chosen the observed proportion of choices in the population of the different strategies, multiplied by 7000 and thus they seem to be much stronger than suggested by the data of most other experiments (Roth and Erev (1995), Camerer and Ho (1999), Nagel and Tang (1999)). Lower initial attractions did not capture at all the slow convergence to A, given signals Y and Z. The reaction to payoffs (of actually played strategies, and also to those of strategies not played in the past), expressed by λ , is also stronger than previous work suggests.

Figure 4a plots the percentage of B action for all possible signals over time (aggregated 5-period blocks). This model captures the spirit of the data (see figure 4b of the actual data also in 5 period blocks): the worse the signal the more likely the B-choice and the slower the convergence to the A-choice.

and Erev 1995). The latter model, in turn, is related to the replicator dynamics of evolutionary game theory (Taylor and Jonker 1978), as Börgers and Sarin (1997) have shown.

We do not claim that our benchmark model is the one that fits best the data, as this is not the purpose of this exercise¹⁰. We are just showing that there exist some models that can fit well the short-run properties of the data. This exercise would not be justified if many other models would also do the job, especially if simpler models were able to capture those properties. Furthermore, at the end we want to use our simulation results to interpret the behavior in our games. With Figure 4c we show that simpler models cannot in fact mimic important features of behavior. There we show the percentage of subjects choosing the *B action under signal Z* for other models. The Z-RE line represents the behavior for standard reinforcement learning, that is, when $\rho = 0$, $\phi = 1$, $\delta = 0$, and $\lambda = 1$, and as initial attractions 7000 as above. Convergence there is too slow, compared to the one observed in the experiment, as one can readily see.

But even leaving all parameters as in the benchmark simulation and then changing only one parameter at a time, produces simulated behavior that is not like the observed one. We set first $\phi = 0$, then $\delta = 0$, and then $\lambda = 1$. These simulations of the relative frequencies of B-action if signal is Z are represented by the Z- $\phi 0$, Z- $\delta 0$, Z- $\lambda 1$ lines, respectively. As with reinforcement learning, convergence is too slow. An opposite problem, that is, a too fast convergence, occurs if we try to do the simulations giving less weight to initial attractions/beliefs. If we multiplied the initial weights by 70 rather than 700, we obtain line Z-70.

We interpret the high weight of initial attractions as an indication that subjects, after reading the instructions, have a definite belief about how one should behave in this game. That is they start with B for low signals (Z, Y and maybe W) and need quite some evidence and time to change behavior. But once they have this evidence, they change faster than standard fictitious play would suggest, as they give significant weight to information of strategies not played previously ($\delta \neq 0$), and they react to payoffs in strong way ($\lambda > 1$).

6. DISCUSSION

In this experiment we have shown that the theoretical predictions of the model in Carlsson and Van Damme (1993) are confirmed by the behavior of real agents in the lab, as long as the experimental subjects have enough time to learn. It is interesting to note that Rubinstein (1989) was written precisely to make evident that the common knowledge predictions may not be robust

¹⁰ Estimating the parameters for this model with experimental data is a delicate question, as Cabrales and García-Fontes (2000) show.

when they required too much rationality from the agents¹¹. That paper actually argued that the theoretical prediction for the e-mail game would not hold in real-life.

Our finding might be interpreted as a validation of the approach used in the applied papers that we cited in the introduction. This should be taken with caution, though. We also find that the complete information version of the game shows a strong tendency to converge to the risk-dominant equilibrium, and convergence is not even faster in the environment with incomplete information. It is thus unclear whether the driving force is, in fact, incomplete information or risk dominance. If it is the latter, the additional modeling complexity incurred by introducing incomplete information may be unnecessary. Incomplete information makes the problem look theoretically cleaner, but behavior may actually not be different in complete and incomplete information games. This is also confirmed in experiments by Heinemann, Nagel and Ockenfels (2002) who depart from a model by Obstfeld (1996) and Morris and Shin (1998).¹²

An additional note of caution should be added, since the comparison of the behavior in our complete and incomplete information games is far from straightforward. One could argue that the behavior under uncertainty is only comparable with the certainty case when the signal is the lowest one. So 15 periods with certainty is comparable (on average) to uncertainty with just 3 periods (since we have 5 signals and each subject experiences only one signal in a period). This, of course, assumes that, under uncertainty, people do not learn anything useful for the behavior with one signal when they receive information about another. Although this does not seem plausible, in order to see if it can make a difference, we have done simulations with the same learning model we used in the previous section for a game in which agents received signals which were completely informative. That is, $\text{signal} = X = \text{true state}$, for X distributed just like in our incomplete information game. The learning parameters are $\rho = 0$, $\phi = 0.9$, $\delta = 0.5$, and $\lambda = 1.5$. Initial attractions are also the initial probabilities multiplied by 7000. Graphical inspection suggests that there is very little difference between the convergence speeds under complete and incomplete information.

One can interpret this lack of difference as saying that the “controllable” part of incomplete information does not add much to the results. There may be a more fundamental lack of common knowledge about preferences, which despite our efforts cannot really be controlled

¹¹ Rubinstein (1989) asserts: “The sharp contrast between our intuition and the game-theoretic analysis is what makes this example paradoxical. The example joins a long list of games such as the finitely repeated Prisoner’s Dilemma, the chain store paradox, and Rosenthal’s game, in which it seems that the source of the discrepancy is rooted in the fact that in our formal analysis we use mathematical induction, while human beings do not use mathematical induction when reasoning.

¹² Although the authors detect a small but significant difference in behavior in complete and incomplete information games, the equilibrium of the incomplete information game better describes the behavior.

(degree of risk aversion, preference for “fairness” and so on). And even with common knowledge of preferences, there could still be lack of common knowledge about the equilibrium, which is why Harsanyi and Selten (1988) introduced risk-dominance. In their words: “risk dominance is important only in those situations where the players would be initially uncertain whether the other players would choose one or the other equilibrium.”

REFERENCES

- Börgers, T. and R. Sarin (1997), "Learning through Reinforcement and Replicator Dynamics," *Journal of Economic Theory*, 77: 1-14.
- Bryant, J. (1983), "A Simple Rational-Expectations Keynes-Type Model," *Quarterly Journal of Economics*, 98:525-529.
- Burzdy, K., D.M. Frankel and A. Pauzner (2001), "Fast Equilibrium Selection by Rational Players Living in a Changing World," *Econometrica* 69:163-189.
- Cabrales, A. and J. Sobel (1992), "On the Limit Points of Discrete Selection Dynamics," *Journal of Economic Theory*, 57:407-419.
- Cabrales, A., W. García-Fontes and M. Motta (2000), "Risk Dominance Selects the Leader: An Experimental Analysis," *International Journal of Industrial Organization*, 18:137-162.
- Camerer, C.F. and T.H. Ho (1999), "Experience-Weighted Attraction Learning in Games: A Unifying Approach," *Econometrica* 67:827-874.
- Charness, G. (1998), "Pre-Play Communication and Credibility: A Test of Aumann's Conjecture, Universitat Pompeu Fabra, Working Paper 293.
- Charness, G and B. Grosskopf (2000), "Cheap Talk and Information," mimeo UPF.
- Carlsson, H. and E.E. van Damme (1993), "Global Games and Equilibrium Selection," *Econometrica* 61:989-1018.
- Cooper, R. and A. John (1988), "Coordinating Coordination Failures in Keynesian Models," *Quarterly Journal of Economics*, 103:441-463.
- Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001), "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica* 69: 1193-1235 .
- Costa-Gomes, M. and V.P. Crawford (2002), "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," manuscript in preparation, University of California, San Diego, and University of York.
- Diamond, P. (1982), "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90:881-894.
- Diamond, D.W. and P.H. Dybvig (1983), "Bank Runs, Deposit Insurance and Liquidity," *Journal of Political Economy* 91:401-419.
- Frankel, D.M., and A. Pauzner (2000), "Resolving Indeterminacy in Dynamic Settings: the Role of Shocks," *Quarterly Journal of Economics*, 115:285-304.
- Friedman, D. (1996), "Equilibrium in Evolutionary Games: Some Experimental Results," *Economic Journal*, 106, 1-25.

- Goldstein, I. (1999), "Interdependent Banking and Currency Crises in a Model of Self-Fulfilling Beliefs," mimeo, Tel-Aviv University.
- Godstein, I, and A. Pauzner (1999), "Endogenous Probability of Bank Runs in a Rational Expectations Model," mimeo, Tel-Aviv University.
- Harsanyi and Selten (1988), *A General Theory of Equilibrium Selection in Games*, Cambridge MA: The MIT Press.
- Hart, O. (1982), "A Model of Imperfect Competition with Keynesian Features," *Quarterly Journal of Economics*, 97:109-138.
- Heinemann, F., R. Nagel, and R. Nagel (2002). Speculative Attacks and Financial Architecture: Experimental Analysis of Coordination Games with Public and Private Information," mimeo, London School of Economics
- Heinemann, F. and G. Illing (1999), "Speculative Attacks: Unique Sunspot Equilibrium and Transparency," mimeo, University of Frankfurt.
- Heinemann, F. (1999), "Unique Equilibrium in a Model of Self-Fulfilling Speculative Attacks: Comment," *American Economic Review*, forthcoming.
- Kajii, A. and S. Morris (1997), "The Robustness of Equilibria to Incomplete Information," *Econometrica* 65:1283-1310.
- Milgrom, P.R and J. Roberts (1991), "Adaptive and Sophisticated Learning in Normal Form Games," *Games and Economic Behavior*, 3:82-100.
- Morris, S., R. Rob and H.S. Shin (1995), " p -Dominance and Belief Potential," *Econometrica* 63:145-157.
- Morris, S., A. Postlewaite and H.S. Shin (1995), "Depth of Knowledge and the Effect of Higher Order Uncertainty," *Economic Theory* 6:453-467.
- Nachbar, J H. (1990), "'Evolutionary" Selection Dynamics in Games: Convergence and Limit Properties," *International Journal of Game Theory*, 19:59-89.
- Nagel, R. (1995), "Unraveling in Guessing Games: An Experimental Study," *American Economic Review* 85:1313-1326.
- Obstfeld, M. (1996), "Models of Currency Crisis with Self-fulfilling Features," *European Economic Review* 40, 1037-1047.
- Rubinstein, A. (1989), "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge,'" *American Economic Review* 79:385-391.
- Sefton, M. and A. Yavaş (1996), "Abreu-Matsushima Mechanisms: Experimental Evidence," *Games and Economic Behavior*, **16**, 280-302.

Shin, H.S. (1996), "Comparing the Robustness of Trading Systems to Higher-Order Uncertainty," *Review of Economic Studies* 63:39-59.

Sonsino, D., I. Erev and S. Gilat (1999), "On Rationality, Learning and Zero-Sum Betting – An Experimental Study of the No-Betting Conjecture," mimeo, Technion University.

Søvik, Y. (2000), "Strength of Dominance and Depths of Reasoning," mimeo, University of Oslo.

Stahl, D.O. and P.W. Wilson (1995), "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior* 10:218-54.

Van Huyck, J.B., R.C. Battalio and R.O. Beil (1990), "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80:234-248.

Van Huyck, J.B., R.C. Battalio and R.O. Beil (1991), "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games," *Quarterly Journal of Economics*, 106:885-910.

Appendix A:

Proposition: The only strategy in this game that survives the iterated deletion of strictly dominated strategies is to play action A irrespective of the signal received.

Proof: Let S_i be the signal received by player $i \in \{1,2\}$, and let $U_i(C | S_i)$ be the expected payoff for player i when using action $C \in \{A, B\}$ after receiving signal S_i .

1. First step of elimination of dominated strategy B: The payoff of action A, given signal T , is calculated by (using the conditional probabilities from table 1):

$$U_i(A | T) = P(X = 90 | S_i = T) \cdot 90 + P(X = 80 | S_i = T) \cdot 80 = \frac{3}{5} \cdot 90 + \frac{2}{5} \cdot 80 = 86, \quad \text{which}$$

is greater than the maximal achievable payoff of 80 when playing action B. Similarly,

$$U_i(A | U) = P(X = 90 | S_i = U) \cdot 90 + P(X = 80 | S_i = U) \cdot 80 + P(X = 70 | S_i = U) \cdot 70 = \frac{3}{7} \cdot 90 + \frac{2}{7} \cdot 80 + \frac{2}{7} \cdot 70 = 81.42857 > U_i(B | U)$$

So playing A is strictly dominant when receiving signal T and U .

2. Second step of elimination: Since playing A is strictly dominant when receiving signals T and U , a player i with signal W who plays B knows that her payoff will certainly be 0 against any rational opponent with signals T and U , so letting $j \neq i$,

$$U_i(B | W) \leq P(S_j \in \{T, U\} | S_i = W) \cdot 0 + (1 - P(S_j \in \{T, U\} | S_i = W)) \cdot 80, \quad \text{where}$$

$$\begin{aligned} & P(S_j \in \{T, U\} | S_i = W) \\ &= P(S_j \in \{T, U\} | X = 80) \cdot P(X = 80 | S_i = W) + P(S_j = U | X = 70) \cdot P(X = 70 | S_i = W) \\ &= \frac{2}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{3} = \frac{1}{3}. \end{aligned}$$

$$\text{Therefore } U_i(B | W) \leq \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 80 < 70.$$

Since,

$$U_i(A | W) = P(X = 80 | S_i = W) \cdot 80 + P(X = 70 | S_i = W) \cdot 70 + P(X = 60 | S_i = W) \cdot 60 = \frac{1}{3} \cdot 80 + \frac{1}{3} \cdot 70 + \frac{1}{3} \cdot 60 = 70,$$

playing A when receiving the signal W is strictly dominant.

3. Third step of elimination: Since players with signal T , U , or W will never play A (according to step 1 and 2), a player with signal Y , who plays B , knows that her payoff will certainly be 0 against an opponent with signal T , U and W , so letting $j \neq i$,

$$U_i(B | Y) \leq P(S_j \in \{U, W\} | S_i = Y) \cdot 0 + (1 - P(S_j \in \{U, W\} | S_i = Y)) \cdot 80, \text{ where}$$

$$\begin{aligned} & P(S_j \in \{U, W\} | S_i = Y) \\ &= P(S_j \in \{U, W\} | Y = 70) \cdot P(X = 70 | S_i = Y) + P(S_j = W | X = 60) \cdot P(X = 60 | S_i = Y) \\ &= \frac{2}{3} \cdot \frac{2}{7} + \frac{1}{3} \cdot \frac{2}{7} = \frac{2}{7}. \end{aligned}$$

$$\text{Therefore } U_i(B | Y) \leq \frac{2}{7} \cdot 0 + \frac{5}{7} \cdot 80 = \frac{400}{7}.$$

Since,

$$\begin{aligned} U_i(A | Y) &= P(X = 70 | S_i = Y) \cdot 70 + P(X = 60 | S_i = Y) \cdot 60 + P(X = 50 | S_i = Y) \cdot 50 \\ &= \frac{2}{7} \cdot 70 + \frac{2}{7} \cdot 60 + \frac{3}{7} \cdot 50 = \frac{410}{7}, \end{aligned}$$

playing A when receiving the signal Y is strictly dominant.

3. Fourth step of elimination: Since a player with signal T , U , W , or Y will never play B (according to step 1 to 3) a player with signal Z , who plays B knows that her payoff will certainly be 0 against an opponent with signals T , U , W , or Y , so letting $j \neq i$,

$$U_i(B | Z) \leq P(S_j \in \{W, Y\} | S_i = Z) \cdot 0 + (1 - P(S_j \in \{W, Y\} | S_i = Z)) \cdot 80, \text{ where}$$

$$\begin{aligned} & P(S_j \in \{W, Y\} | S_i = Z) \\ &= P(S_j \in \{W, Y\} | X = 60) \cdot P(X = 60 | S_i = Z) + P(S_j = Y | X = 50) \cdot P(X = 50 | S_i = Z) \\ &= \frac{2}{3} \cdot \frac{2}{5} + \frac{1}{2} \cdot \frac{3}{5} = \frac{17}{30} \end{aligned}$$

Therefore $U_i(B | W) \leq \frac{17}{30} \cdot 0 + \frac{13}{30} \cdot 80 = \frac{104}{3}$.

Since, $U_i(A | Z) = P(X = 60 | S_i = Z) \cdot 60 + P(X = 50 | S_i = Z) \cdot 50 = \frac{2}{5} \cdot 60 + \frac{3}{5} \cdot 50 = \frac{270}{5}$,

then playing A when receiving the signal Y is strictly dominant. Thus, a player with any signal S will always choose A. *Q.E.D.*

Appendix B:

Instructions: Uncertainty treatment (translated from Spanish into English)

Thank you for participating in this experiment, a project of economic investigation. Your earning depends on your decisions and the decisions of the other participants. There is a show up fee of 500 pesetas assured. From now on until the end of the experiment you are not allowed to communicate with each other. If you have some question raise your the hand and one of the instructors will answer the question in private. Please, do not ask aloud. Thank you very much.

The rules are equal for all the participants. The experiment consists of _____ rounds. In each round you are matched with a different participant, chosen at random. You will never know who your partner is nor your partner will know who you are. The rules are the same in all the rounds.

In each round you will have to make a decision. In each round, as you will see below, there is a hidden value, generated randomly in each round. You will receive a hint, useful to know something on that hidden value and to make a good decision. Your partner also will receive a hint on the same hidden value, but that does not mean that this hint must be the same one as yours. What you win (and what your partner wins) depends on the hidden value, your decision and the decision of your partner.

In each round, you will be able to choose between the decision A or decision B. The following matrix describes the associated payments to each decision:

		<i>Decision of your partner</i>	
		<i>A</i>	<i>B</i>
<u>Your decision</u>	<u>A</u>	X <i>X</i>	X <i>0</i>
	<u>B</u>	0 <i>X</i>	80 <i>80</i>

The matrix is the same one for all the participants. The rows indicate your decision **A** or **B**, and the columns the decision of your partner (also *A* or *B*). In each cell the value in the left upper angle is the payment which you receive (in bold), and the one which is in the right lower angle, the one your partner receives (italics).

If you play B and your partner also plays B, both of you receives 80 Pesetas. If you play A and your partner play B, you receive X points and your partner receives 0 points. In order to see if you have understood the matrix well, ask yourself what happens if you both A. Answer, you both receive X. Neither you nor your partner will know the exact value of X when making a decision. If you have some doubt, or you are not familiarized with this type of matrices, please raise your hand and an instructor will answer your question.

The X that appears in the matrix of payments is the hidden value of the game. This X is extracted out of five possible values, randomly and independently drawn for each round. That is, the value X of the previous round does not influence at all the value X of the following round. The five possible values of X are 90, 80, 70, 60, 50. The five values are equally probable: there is a 20% chance that the value is 90, a 20% chance that it is 80, a 20% chance that it is 70, a 20% chance that it is 60, and a 20% chance that it is 50. The value of X will be the same one for you and your partner.

To help you to make a good decision, you will receive a hint of value of X. With this, you will be able to sharpen your knowledge of X, although you will not know it exactly when making your decision. The hint of your partner is on the same X, but it is independent of your hint, it can be the same one or it can be different. The following table explains how the hint and the true value X are related.

Value of X	90	80	70	60	50
Hint	T or U	T or U or W	U or W or Y	W or Y or Z	Y or Z

Explanation of the table:

- Hint T: X can be 80 or 90
- Hint U: X can be 90, 80 or 70
- Hint W: X can be _____ (fill in what it can be).
- Hint Y: X can be 70, 60 or 50
- Hint Z: X can be 60 or 50

Given a value of X, each one of the possible hints for that value is equally probable.

Information after each round:

After each round you will be informed of the decision of your partner, your payment and the payment of your partner. Also we will inform you about the value of X.

Summarizing, in each round a different partner, who you will not know and selected at random, will be matched with you. You will receive a hint on the value of X, which is unknown to you and independent of the values of the previous rounds. Your partner also will receive a hint on the same X, but the hint can be different. You will have to make a decision A or B. When all the participants also have made their decision, you will know which was the value of X, the decision of your partner and the payment that you receive in this round.

The gains: At the end of _____ rounds we add your points and we convert a point to 0,5 pesetas, so that 1000 points are 500 Pesetas. In addition you will receive 500 Pesetas.

If you have some question, please raise your hand.

We would be interested in some comments on how you arrived at your decision. Please write them in the comment sheet.

Appendix C: Figures

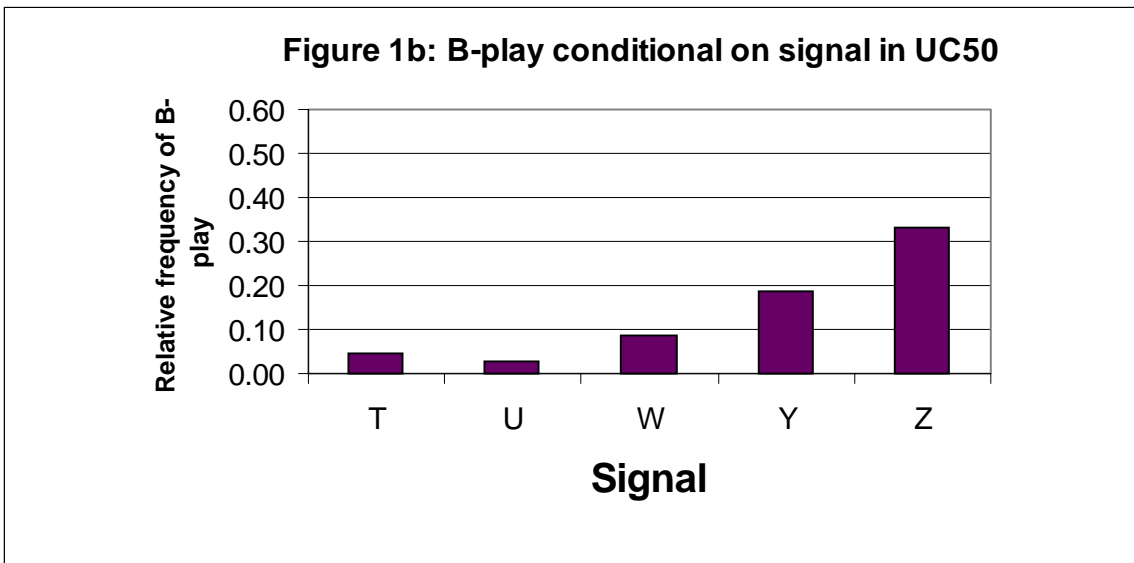
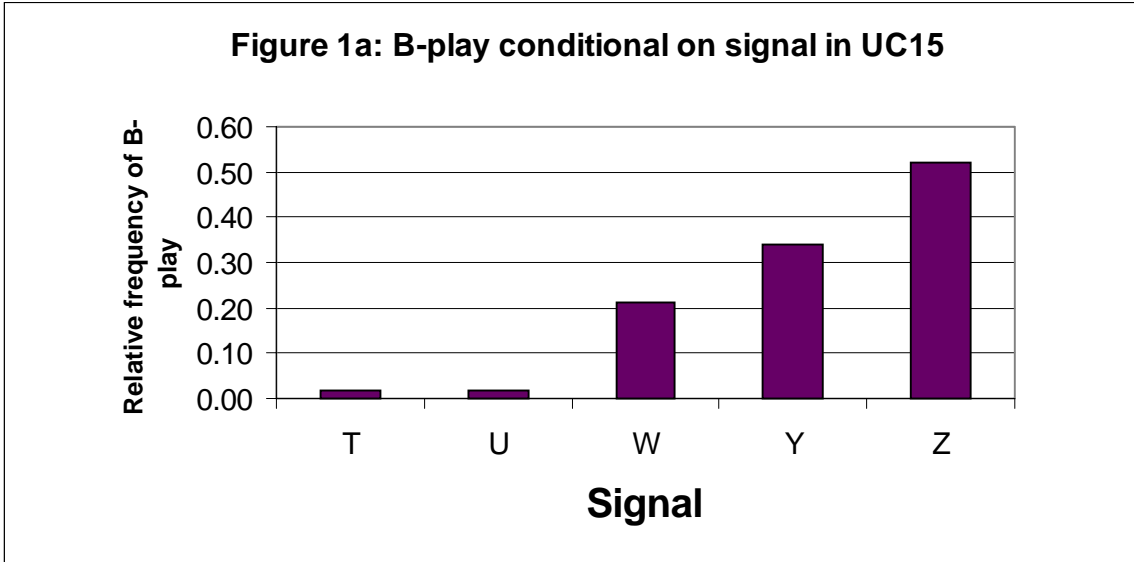


Figure 2a: B-play over time conditional on signal in UC15

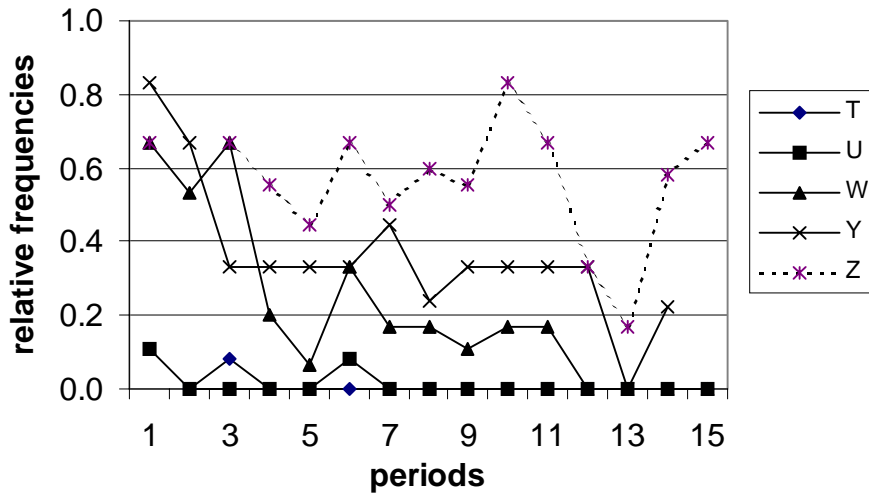


Figure 2b: B-Play over time conditional on signal in UC50

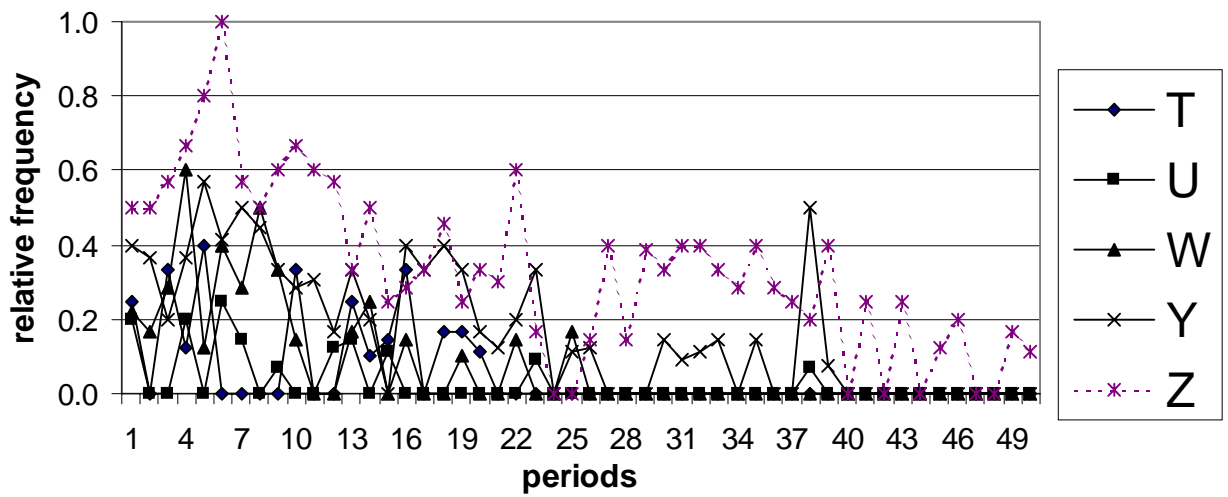


Figure 3: B-play over time in UC15+UC50
when signal Y or Z and B-play in C15

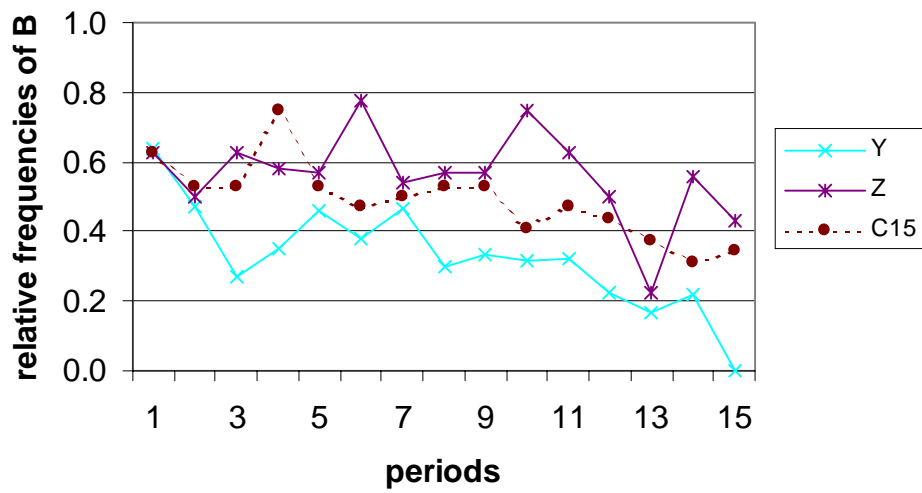


Figure 4a: Simulation of relative frequency of B choices in UC50 conditional on signal

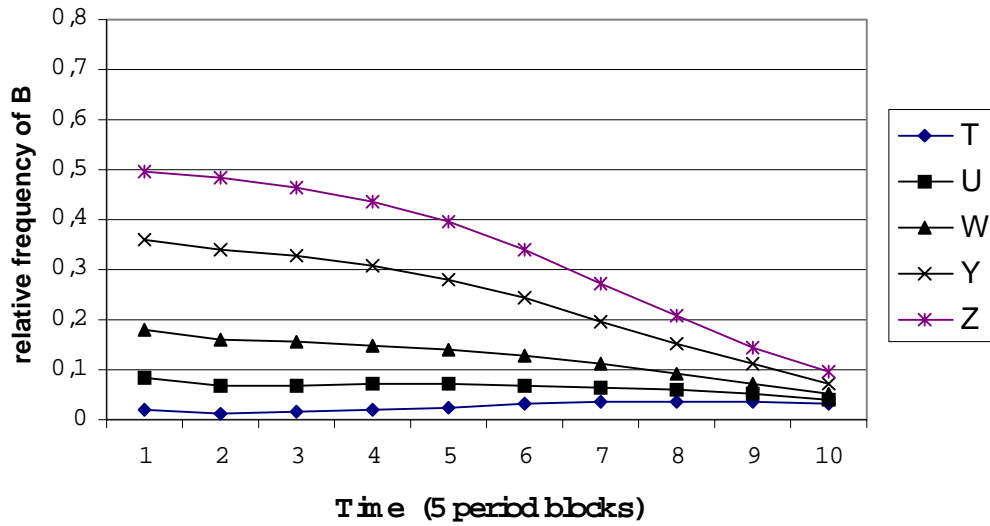


Figure 4b: Relative frequency of actual B-choices in UC50 conditional on signal (in 5 period blocks)

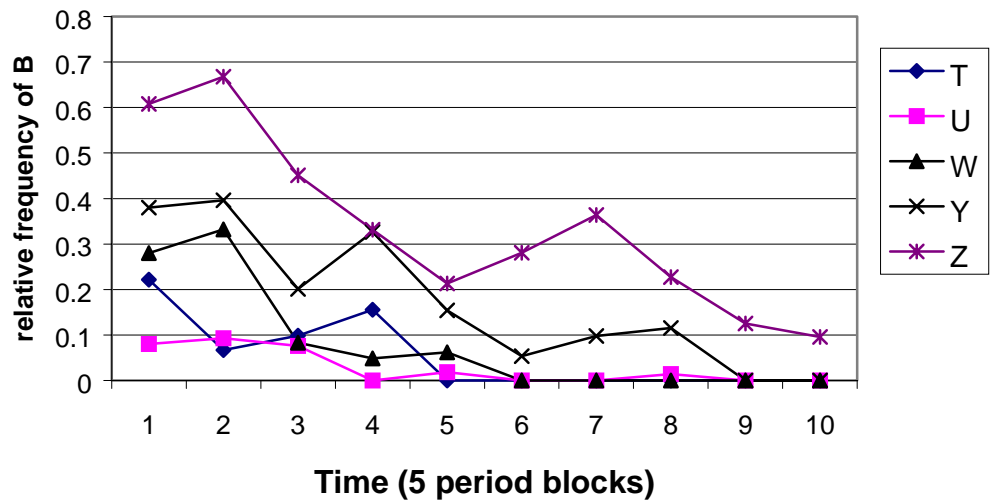


Figure 4C: Different simulations of relative frequency of B choices in UC50 conditional on signal Z

