

Risk dominance selects the leader: An experimental analysis

Antonio Cabrales^{a,*}, Walter García-Fontes^a, Massimo Motta^{a,b}

^a*Department of Economics, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain*

^b*European University Institute, Florence, Italy*

Abstract

We perform an experimental analysis to test whether the risk dominance prediction is supported by the behavior of laboratory agents, in a 2×2 coordination game whose equilibria are not Pareto ranked. This type of game arises very often in studies of industrial organization and international trade, and we extract the parameters for the experiment from a vertical product differentiation model with two asymmetric players choosing first qualities and then prices. We show that the higher the degree of asymmetry of the game, the higher the predictive power of the risk dominance criterion. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Risk dominance; Coordination game; Vertical product differentiation; Asymmetric players

1. Introduction

The game known as ‘Battle of the Sexes’ (Luce and Raiffa, 1957),¹ which we subsequently denote by BoS, is the simplest game with two pure-strategy equilibria where the players have divergent interests over the different equilibria. Most standard criteria of selection among multiple equilibria (such as perfectness, properness, strategic stability and Pareto dominance) have no selective power in

*Corresponding author. Tel.: +34-93-542-2765; fax: +34-93-542-1746.

E-mail address: cabrales@upf.es (A. Cabrales)

¹This game is also called Chicken, as in Binmore (1992).

this game. A notable exception is the criterion of risk dominance proposed by Harsanyi and Selten (1988).

To test the prediction of the risk dominance criterion we have performed experiments on the BoS. The game is shown in Table 1. Matrix 1 is symmetric, while matrices 2–6 show increasing asymmetries in payoffs. The main hypotheses we tested were whether the risk dominant equilibrium is selected by the subjects in this game and whether the degree of asymmetry matters in this selection. We were also interested in how the agents adjust their play towards an equilibrium and whether play converges or not. This could serve to elucidate whether some evolutionary models that have suggested risk dominance as an equilibrium selection criterion are useful in explaining the facts.

The experimental results give support for the use of risk dominance as an equilibrium selection criterion in this type of game. This support has to be qualified, though, in the following sense: what we find is that the actions leading to the risk dominant equilibrium are chosen more often than the alternative actions, but they are not always chosen. Furthermore, the frequency of the risk dominant outcome is much higher than the alternatives when the asymmetry is large but only marginally so as the asymmetry becomes smaller.

We tested different learning protocols and we found that fictitious play does a worse job at explaining the data than unsophisticated stimulus–response models which are related to evolutionary processes. In addition to their intrinsic game theoretic interest, we wanted to evaluate different learning rules in this game because some recent research (Kandori et al., 1993; Young, 1993) has shown that

Table 1
Payoffs of the game

Matrix 1:				Matrix 2:			
		Choice of B				Choice of B	
		0	1			0	1
Choice of A	0	29, 29	594, 365	Choice of A	0	123, 78	600, 364
	1	365, 594	338, 338		1	365, 567	348, 341
Matrix 3:				Matrix 4:			
		Choice of B				Choice of B	
		0	1			0	1
Choice of A	0	193, 114	604, 363	Choice of A	0	248, 142	608, 362
	1	365, 545	356, 343		1	364, 526	362, 345
Matrix 5:				Matrix 6:			
		Choice of B				Choice of B	
		0	1			0	1
Choice of A	0	292, 164	612, 362	Choice of A	0	327, 181	615, 361
	1	364, 509	367, 346		1	364, 496	371, 347

under evolutionary learning processes risk dominant equilibria will be played more often than the other equilibria.

We also find that observed play does not show a strong tendency towards convergence. Although miscoordination is to be expected at the beginning of play (and in the absence of strong focal points), in a coordination game like the one we are analyzing this is a bit surprising when it lasts for 25–30 periods.² But it is less surprising when one considers that the agents' information about history consisted only of the outcomes from the games they had played, and since they were randomly matched each period there was no obvious 'correlation device' they could use, not even history, since different players had different histories which were unknown to each other.

The BoS is important for industrial organization because there are many situations which have multiple equilibria that are ranked differently in the preferences of the agents. In particular, it can be viewed as a simplified version of a vertical product differentiation model where two firms first decide on the quality they want to bring to the market and then choose prices. We performed experiments in a BoS game derived from our vertical product differentiation model. The payoffs in the experiment are those of the equilibrium strategies (suitably renormalized) for the differentiation model. This implies a strong discretization of the strategy space (the continuum in the original differentiation model), which makes the experimental game easy to analyze for the subjects without losing much in terms of the strategic situation.

This vertical product differentiation model endogenously gives rise to a 'leader' and a 'follower', the former being the firm with higher quality and profit at equilibrium (product market competition prevents firms from choosing the same quality in equilibrium). If the firms are different, the magnitude of the initial asymmetry (on whose nature we say more below) is crucial in determining the outcome of the game. When a firm has a strong enough initial advantage over the rival, the only possible equilibrium is one where the leader at the equilibrium will be the firm which enjoys the advantage. However, when initial asymmetries are not too large, the game admits two possible equilibria. In the first the leader is the firm enjoying the initial advantage, and in the second the leader is the firm starting with an initial disadvantage. We show that the concept of risk dominance unambiguously selects the leader. Indeed, the equilibrium where the leader is the firm which enjoys an initial advantage always risk dominates the other equilibrium, independently of the magnitude of the advantage.³

Several interpretations of this model and of the initial asymmetries are possible.

² See, e.g. Friedman (1996).

³ Provided such advantage exists. Risk dominance predicts the mixed strategy equilibrium in the symmetric game

One of them is that the firms differ in their initial quality levels (which we can think of as historically given) and they have to incur some adjustment costs to adjust the quality of their product up to the desired level. Under this interpretation, the model gives some insights about the extent to which a current leader (the firm starting with the higher initial quality) will be able to persist as the industry leader, or whether the initial advantage can be overturned, with the initial follower establishing itself as the top quality firm at the equilibrium.

Viewed in this way our paper is related to the wide literature in industrial organization and international trade which aims at establishing conditions under which asymmetries between firms and between countries tend to widen or shrink over time. This problem has been studied under many points of view, and with different answers. As an example, we can cite patent race models (see, e.g. Reinganum, 1985, in a partial equilibrium model, and Grossman and Helpman, 1991, in a general equilibrium model with innovation and growth), learning-by-doing models (Dasgupta and Stiglitz, 1988) or switching-cost models (Farrell and Shapiro, 1988). A recent treatment in a dynamic model of a duopolistic industry with sufficiently general features has been provided by Budd et al. (1993). Brezis et al. (1993), Flam and Helpman (1987) and Motta et al. (1997) study the problem from an international perspective, and analyze whether latecomers can overtake countries starting from a higher level of development.

The other branch of the literature to which our paper is connected is the experimental literature on coordination games. The closest analog to our experiments we have found are Guyer and Rapoport (1972), Straub (1995) and Friedman (1996). They study coordination games but do not address specifically the issue of risk dominance or the effects of asymmetries on the choice of equilibrium strategies. Other previous experimental research on similar games has tended to focus on symmetric games, for instance Cooper et al. (1989, 1993) or Van Huyck et al. (1990). Charness (1998) and Clark et al. (1999) study a 2×2 pure coordination game where the Pareto ranking is opposite to the risk dominance ranking, while Sefton (1998) analyzes econometrically a learning model based on a coordination game.

The paper is organized as follows: in Section 2 we present the theoretical model with asymmetric players and multiple equilibria, while in Section 3 we discuss the selection concept of risk dominance. In Section 4 we present the experimental design. The results are summarized in Section 5, and some learning models are tested in Section 6. Concluding remarks are presented in Section 7.

2. The model

We consider a version of the vertical product differentiation model in the tradition of Shaked and Sutton (1982) and Mussa and Rosen (1978). We assume a population of consumers who have utility function $U = \theta u - p$ if they buy one unit

of the differentiated good and $U = 0$ if they do not buy. The symbols u and p stand for quality and price of the good, and θ represents a taste parameter. The distribution of θ is assumed to be uniform with $\theta \in [0, \bar{\theta}]$ and a density S .

There exist two firms in the industry, L and H . In the first stage of the game they decide on the quality $u_j \in \Re^+$ they want to produce and incur a fixed (i.e. independent of the quantity produced) cost of quality $F_j = k_j \frac{u_j^2}{2}$, with $j = L, H$. The cost of quality can be thought of as R&D investments or advertising outlays. The quadratic form is a standard assumption which greatly simplifies calculations. Any convex function would give the same qualitative results. Note that the two firms do not necessarily have the same technology: firm L is at least as efficient as firm H , with $1 = k_L \leq k_H = k$. Therefore, the parameter k is a measure of the asymmetry existing between the two firms. We take marginal production costs to be constant and, without loss of generality, we set them equal to zero.

In the second and final stage of the game, firms choose prices. We look for the sub-game perfect equilibrium in pure strategies of this game. For reasons of space, we omit the derivations and just summarize the results.⁴

It turns out that there are two equilibria of this game. In the first one, $E_L = (u_0^*, u_1^*)$, it is the more efficient firm L which produces the top quality (without loss of generality, we denote with 0 a higher quality: $u_0 > u_1$). In the second equilibrium, $E_H = (u_0^{**}, u_1^{**})$, it is the less efficient firm H . In each equilibrium, the top quality firm makes considerably higher profits than the low quality firm. Because firm L is (weakly) more efficient, the top quality levels and the top quality seller's profits are (weakly) higher at equilibrium E_L than at equilibrium E_H . When $k = 1$, the two firms are equally efficient and the two equilibria are identical.

It is possible to show that equilibrium E_L exists for all the values of the parameters, whereas E_H is not always an equilibrium. Indeed, there exist high enough values of parameter k ($k = 1.5894$ is approximately the threshold value above which this equilibrium collapses) for which this equilibrium ceases to exist.⁵

There exist other examples of vertical product differentiation models which share the same basic features as the model analyzed here. Consider for instance the following variation of the model. Firms have exactly the same technologies ($k = 1$ for both), but when the game starts the firms have inherited different levels of quality (which can be interpreted as the consequence of past levels of R&D or

⁴Details are available from the authors upon request. Supplemental materials, such as omitted derivations and the data from the experiments, are posted at the following URL: <http://www.econ.ku.dk/cie/ijio/exptsup.htm>

⁵For an intuitive argument, consider the extreme case where firm H is infinitely inefficient. If k tends to infinity, then firm H will choose at the candidate equilibrium a top quality $u_0^{**} = \epsilon$, with ϵ arbitrarily small and infinitely small profits. Firm L 's profits are even lower, since the bottom quality firm always makes less profit, and therefore at a small cost the L firm can produce a quality higher than ϵ , become the top firm and earn higher profits.

advertising expenditures). In the first period of the game, they can update the quality of the good by incurring some adjustment costs; in the second period, they compete on prices. Two equilibria might arise: one where the firm endowed with the larger initial quality will still be producing the higher quality at the new equilibrium (persistence of dominance) and one where it is the initial laggard firm which provides the high quality (leapfrogging); the latter equilibrium ceases to exist when the difference in initial quality levels is too large.⁶ This model has been analyzed, in an international trade context by Motta et al. (1997).⁷

3. Risk dominance

The game played by two firms in the context of a vertical product differentiation model we described above has two strict Nash equilibria, for a range of values of the parameter k . One of the few solution concepts that selects between equilibria in our game,⁸ is the concept of risk-dominance introduced by Harsanyi and Selten (1988).

In this section we will define the concept of risk-dominance for 2×2 games, which we will use in the experiments. We will show that risk dominance selects the equilibrium where the firm with the low cost of quality is the leader when only two quality levels (the equilibrium ones, as in the experimental design) are possible.

To motivate the concept of risk dominance consider the following game, which has two strict equilibria in pure strategies, (U_1, U_2) and (V_1, V_2) , and its discussion by Harsanyi and Selten (1988, p. 82)

	U_2	V_2
U_1	99,49	0,0
V_1	0,0	1,51

If player 1 expects that player 2 will choose U_2 with probability of more than 0.01 it is better for him to choose U_1 . Only if player 2 chooses V_2 with a probability at least 0.99, player 1's strategy will be the more profitable. In this

⁶Like in the model we present here, it turns out that the 'leapfrogging' equilibrium is always risk-dominated by the 'persistent leader' equilibrium.

⁷See also, Cabrales and Motta (1996) for another model with similar features.

⁸Standard refinements like perfectness, properness, or strategic stability do not select among strict Nash equilibria. Also, in this game there are no symmetric equilibria for $k > 1$ and neither equilibrium Pareto dominates the other (taking into account only the welfare of the players, the firms, and not the consumers).

sense U_1 is much less risky than V_1 . Now let us look at the situation of player 2. His strategy V_2 is the better one if he expects player 1 to select V_1 with a probability of more than 0.49, and U_2 is preferable if he expects U_1 with a probability greater than 0.51. In terms of these numbers V_2 seems slightly less risky than U_2 . It is obvious that player 1's reason to select U_1 rather than V_1 is much stronger than player 2's reason to select V_2 rather than U_2 .

The reasoning is that the equilibrium point (U_1, U_2) involves the less 'risky' choice for 1, and the equilibrium (V_1, V_2) involves the less 'risky' choice for player 2. If the 'riskiness' for the two players pointed to the same equilibrium the reasons for selection would be clearer, but their point is that even when this is not so, one can still select between the equilibria by weighting in some way the two players' risk motives. The concept of risk dominance is a way to make precise this type of reasoning. This criterion compares the product of gains from correct predictions and the equilibrium with the largest product is the one that risk dominates.

Let a 2×2 game have the following payoff matrix:

		Player H	
		H_0	H_1
Player L	L_0	l_{00}, h_{00}	l_{01}, h_{01}
	L_1	l_{10}, h_{10}	l_{11}, h_{11}

where the payoffs are such that $E_L = (L_0, H_1)$ and $E_H = (L_1, H_0)$ are strict Nash equilibria,⁹ and let $GL_0 = l_{01} - l_{11}$. GL_0 is the gain made by player L by correctly predicting that the other player will play as in E_L (and best responding to the prediction) instead of incorrectly predicting that the other player will play as in E_H (and best responding to the prediction). Similarly, let $GH_0 = h_{01} - h_{00}$. $GL_1 = l_{10} - l_{00}$. $GH_1 = h_{10} - h_{11}$. We say that equilibrium E_L risk dominates equilibrium E_H when $GL_0 \cdot GH_0 > GL_1 \cdot GH_1$.

Harsanyi and Selten provide an axiomatic justification for the criterion. Risk dominance is the only equilibrium selection criterion that is not affected by a relabeling of strategies and players, nor by a change of the payoffs that maintains the structure of the best response correspondence, nor by strengthening the payoffs of the selected equilibrium.

Besides the intuition and the axiomatization provided by Harsanyi and Selten, there are more theoretical and empirical reasons why risk dominance could be

⁹We change the order of the strategies with respect to the example in the previous paragraph for consistency with the game of quality choice, where the strategies have a natural ordering in terms of quality levels.

considered a good equilibrium selection criterion. In evolutionary models, like Kandori et al. (1993), Young (1993) and Fudenberg and Harris (1992) with noise due to random shocks to payoffs, or to mutation/experimentation by new or uninformed players, observed play follows a stochastic process. The limiting distribution of that stochastic process when the noise is small puts most of the weight in the state where all agents are playing according to the risk dominant equilibrium. Another theoretical argument in favor of risk dominance arises in games with some uncertainty about the payoffs. Suppose players are not completely sure whether the game played is actually one with two strict equilibria (although this is the most likely event) or one where either U or V are strictly dominant. The only equilibria of the game with that type of uncertainty is such that both players play the strategies from the risk dominant equilibrium in the game without uncertainty (Carlsson and Van Damme, 1993). Perhaps the most appealing argument for risk dominance is that in the experimental evidence available so far, as discussed in the Introduction, the risk dominant equilibrium is chosen by the subjects.

Let us now apply the risk dominance criterion to our model. Recall that, as discussed in Section 2, L is the low cost firm and H is the high cost firm, and that E_L is the equilibrium where the low cost firm is the leader while E_H is the equilibrium where the high cost firm is the leader. The strategies of our model are the low or high qualities corresponding to the two possible equilibria, that is $E_L = (u_0^*, u_1^*)$ and $E_H = (u_0^{**}, u_1^{**})$. Therefore the strategies of the 2×2 game are: $L_0 = u_0^*$, $L_1 = u_1^{**}$, $H_0 = u_0^{**}$ and $H_1 = u_1^*$.

The payoffs of the 2×2 matrix are $l_{00} = \pi_L(u_0^*, u_0^{**})$, $l_{01} = \pi_L(u_0^*, u_1^*)$, $l_{10} = \pi_L(u_1^{**}, u_0^{**})$, $l_{11} = \pi_L(u_1^{**}, u_1^*)$, $h_{00} = \pi_H(u_0^*, u_0^{**})$, $h_{01} = \pi_H(u_0^*, u_1^*)$, $h_{10} = \pi_H(u_1^{**}, u_0^{**})$ and $h_{11} = \pi_H(u_1^{**}, u_1^*)$, where $\pi_i(u_L, u_H)$ with $i = L, H$, is the profit level that firm i attains when the L firm chooses u_L and the H firm chooses u_H .

In our case, GL_0 is what player L gains by correctly forecasting that the other player will play the equilibrium where L itself is the leader, instead of forecasting wrongly that the other equilibrium holds. GL_1 represents the gains for player L of forecasting rightly that the other player will play the equilibrium where H is the leader. The interpretation of GH_0 and GH_1 is analogous.

It is easy to show that $GL_0 \geq GH_1$ and $GH_0 \geq GL_1$, and the equality only holds when $k = 1$, that is, when both firms have identical costs of quality. Thus, for the game we are studying, the risk dominance criterion selects the equilibrium where the leader is the low cost firm.

Equilibrium selection can be interpreted in a strong way and in a weaker way. The strong interpretation is that we will never (or rarely) observe players choosing the strategies that lead to an equilibrium that is not selected. The weaker interpretation is that the likelihood of observing players using the strategies that lead to each equilibrium is related to the ‘degree’ of risk dominance. To make this more precise we propose the measure of this degree to be $\frac{GL_0 \cdot GH_0}{GL_1 \cdot GH_1 + GL_0 \cdot LH_0}$. Notice that this ratio is invariant to affine transformations of the payoffs, so it can

be used to measure the strength or degree of risk dominance. The larger this ratio, the more risk dominated E_H is, and so the likelier it is that we will observe players using the E_L strategies. We feel that the weaker hypothesis is more reasonable, not only because it is more difficult to reject, but also because most of the reasons in favor of risk dominance involve uncertainty or bounded rationality (see, Harsanyi and Selten, 1988, p. 89), and in these circumstances it would be hard to expect the strong hypothesis to be satisfied. Our interpretation is supported by the experimental data we describe below.

4. Experimental design

Even though in the motivating model each agent has a continuum of strategies, in the experiments we chose to restrict the number of strategies to two. There are several reasons for this. Firstly, it can be shown that risk dominance selects the same equilibrium in the game with 2 strategies per player as well as in a game with a much less coarse discretization of the strategy space.¹⁰ Besides, by reducing the dimensionality in this way we partly compensate for the problems created by the smaller incentives that agents have in laboratory experiments (computational complexity would justify a more careful behavior in the real world, where both time available and the payoffs are likely to be higher). It is also unclear whether the smaller strategy space entails a departure from realism. Using a continuum of levels of quality allows us to use calculus tools which greatly simplify the theoretical analysis, but it may be that two quality levels comes closer to the available number of strategies in the real world case.

The game played in the experiment was a 2×2 bimatrix game. Each agent chose one of two actions; 0 or 1. There were two types of players; A and B. The role of the subjects was randomly assigned to them at the beginning of the session and was not changed for the duration of the session. The game was played for 25 or 30 periods depending on the session. Each player was randomly matched with an anonymous opponent after every period. The agents were informed of the number of times the game was played and they had two initial practice rounds for which they received no payment. The payoffs were given by payoff matrices in Table 1, which were known to all the players, so the game was one of complete information. These matrices are derived from the model we analyzed in Section 2, as we specify below.

The experiments were conducted using 6 or 7 subjects of each type, depending on the session. The subjects were undergraduate students in the Faculties of Economics and Business and of Humanities at Universitat Pompeu Fabra in Barcelona. Players were seated at separate computer terminals and given a copy of

¹⁰For details see footnote 4.

the instructions which can be found in Appendix A translated to English (the originals were in Spanish). They could ask questions if they did not understand the instructions, but there were essentially no questions, other than about how to read the payoff matrix.

To induce the payoffs in terms of utility and control for risk aversion we used the binary lottery procedure (see, Roth and Malouf, 1979). Each player received the payoff in points which then determined the probability of receiving a monetary prize. At the end of each period we conducted a lottery in which winning players received a prespecified prize, and the losers received nothing. The probability of winning was given by the number of points divided by 1000. This procedure guarantees that expected utility maximizers will try to maximize the number of points independently of their attitudes towards risk.¹¹

We resorted to this procedure because risk dominance is a concept which is not invariant to risk aversion. It is possible to find examples of risk preferences in our game such that if you write the matrix for the game in terms of monetary payoffs (that is, if agents are risk neutral) one equilibrium dominates but for other risk preferences the risk dominant equilibrium is a different one. This led us to control for risk aversion, even though some authors question the general validity of the procedure (see, Selten et al., 1995, and the remarks on the use of the binary lottery by Roth in the introduction to Kagel and Roth, 1995).

To see if the use of this procedure could bias the results we ran some control sessions (7 and 8) where points were converted into pesetas (the Spanish currency) at a fixed deterministic exchange rate, which was known the experimental subjects. There were no major qualitative changes in behavior in those control sessions, so at least in our design this issue seems to be relatively unimportant.

Sessions 2, 4 and 6 were run right after 1, 3 and 5, respectively, and the same subjects were used. A new matrix was drawn and the subjects' roles were drawn randomly again. They were asked to play another 25 periods.¹² This was done to test for experience effects with a different matrix. Indeed it turns out that previous experience made players coordinate more easily in the new game, even though subjects were assigned different roles than in the previous session and a different matrix was chosen to play.

The agents saw on their screens the history of actions and outcomes of the matches they had been in, which was automatically updated after every period. They knew what they had played and what their different opponents had played in

¹¹Notice that expected utility is invariant with respect to affine transformations, so by normalizing the utility of the prize to 1000 and the utility of no prize to 0 the expected utility is equal to the expected number of points.

¹²When subjects played only one session, this lasted 30 periods. When subjects played two subsequent sessions, each of them lasted 25 periods. This was done to shorten the time they had to devote to the experiment.

every period, and how much they had earned in every period, both in terms of points and in terms of pesetas, as well as their cumulated pesetas. But they did not know what the history of their opponents was, nor had they any aggregate statistic of how many players had chosen every strategy in every period. We chose this design to avoid that expectations were too easily coordinated, which might decrease the informativeness of the individuals' actions.

The payoffs of the game, shown on Table 1, correspond to the payoffs for equilibrium strategies in the vertical product differentiation game described in Section 2, after a renormalization to make the numbers fall between 0 and 1000. This renormalization makes it easier for the subjects to understand the conversion of points into probabilities in the binary lottery procedure. We used six matrices, numbered 1–6, which correspond respectively to the values of $k \in \{1, 1.1, 1.2, 1.3, 1.4, 1.5\}$. Recall that for $k = 1$ the game is symmetric and that the asymmetries rise with the value of k . For all matrices, strategy 0 is the level of quality chosen by a firm in the equilibrium in which that firm produces the higher quality and strategy 1 is the level of quality in the equilibrium in which that firm produces the lower quality. The payoffs in the matrix are the renormalized payoffs corresponding to the payoffs each firm would earn in the model described in Section 2. For instance, in matrix 3 the pair (0,0) would represent the situation where player L is playing the quality level 25.28 and R is playing 23.08 and $k = 1.2$, and the payoffs are the ones that correspond to this strategy profile (which is not an equilibrium). Notice that $k < 1.56$ so we are considering only situations where both equilibria exist.¹³

Table 2 summarizes session numbers, payoff matrices used, whether the session had experienced players (in brackets the session where the same subjects had played) and whether the session used the binary lottery procedure or not.

5. The data and results

Table 3 and Fig. 1 summarize the results obtained with the pooled data from the 12 experiments.¹⁴ In Figs. 1 and 2 we compute a cubic regression curve with 95% confidence intervals for each type of equilibrium, in order to identify graphically the incidence of the two types of equilibrium. Graphically, it can be seen that equilibrium (0,1) was played more often than (1,0), thus the risk dominance criterion makes the right prediction more often. We also computed a sign rank statistic to test whether the difference between the number of equilibria (0,1)

¹³The subjects were not informed about the economic interpretation in terms of a quality choice game, so as not to bias their perception in terms of the 'prestige' possibly associated to producing a high quality.

¹⁴Fig. 1 also includes the two trial periods performed in each session.

Table 2
Characteristics of each session

Number	Matrix	Experience (and previous session when experienced)	RM
1	5	No	Yes
2	2	Yes, 1	Yes
3	6	No	Yes
4	3	Yes, 3	Yes
5	1	No	Yes
6	4	Yes, 5	Yes
7	5	No	No
8	2	No	No
9	6	No	Yes
10	5	No	Yes
11	2	No	Yes
12	3	No	Yes
13	4	No	Yes

minus the number of equilibria (1,0) is greater than 0.¹⁵ We pooled the observations for groups of five periods. The results are reported in Table 4, also providing evidence that the number of risk dominant equilibria is larger.


As we noted in Section 3 the degree of risk dominance can be measured by the

Table 3
Outcomes of the sessions

Matrix	% of (0,1)	% of (1,0)	% of (0,0)	% of (1,1)
First period				
All periods	35.6	17.2	23.0	24.1
1	0.0	16.7	0.0	83.3
2	15.0	40.0	25.0	20.0
3	15.0	40.0	25.0	20.0
4	57.1	21.4	0.0	21.4
5	30.0	15.0	20.0	35.0
6	71.4	0.0	28.6	0.0
All sessions				
All matrices	41.9	17.2	28.3	12.6
1	14.7	42.7	18.0	24.7
2	26.8	28.3	31.7	13.4
3	40.0	11.1	32.5	16.4
4	41.9	13.1	33.9	11.1
5	44.8	16.3	26.9	12.0
6	72.5	1.8	20.5	5.2

¹⁵The details for this and all subsequent tests are shown in Appendix B.

Difference between equilibria: all experiments



Note: 95 % confidence intervals around the mean (---)

Fig. 1. Difference between equilibria: all experiments.

ratio $\frac{GL_0 \cdot GH_0}{GL_1 \cdot GH_1 + GL_0 \cdot GH_0}$, which is an increasing function of k . It is clear that the proportion of times for which the risk dominant equilibrium is played increases with k as can be seen in Fig. 2, confirming the weaker interpretation of risk dominance we proposed.¹⁶ We confirmed this graphic visualization by running a t -test of the hypothesis that the difference in the number of equilibria is linearly increasing with k , which is an increasing function of the degree of risk dominance of the equilibrium (1,0) of the matrices. We also divided the periods in sequences of 5. The results are reported in Table 5. The results strongly show that we cannot reject the hypothesis that the difference is linearly increasing.

Fig. 2 shows that as time evolves, agents tend to play the risk dominant equilibrium more often, in the spirit of the theories of learning. The next section

¹⁶ We tested the robustness of this result by replicating the experiments for matrices 5 and 6, by using a direct compensation scheme instead of the binary lottery procedure and by having agents play with a different matrix after gaining experience with another matrix. This result and the ones reported below do not change. For details see footnote 4.

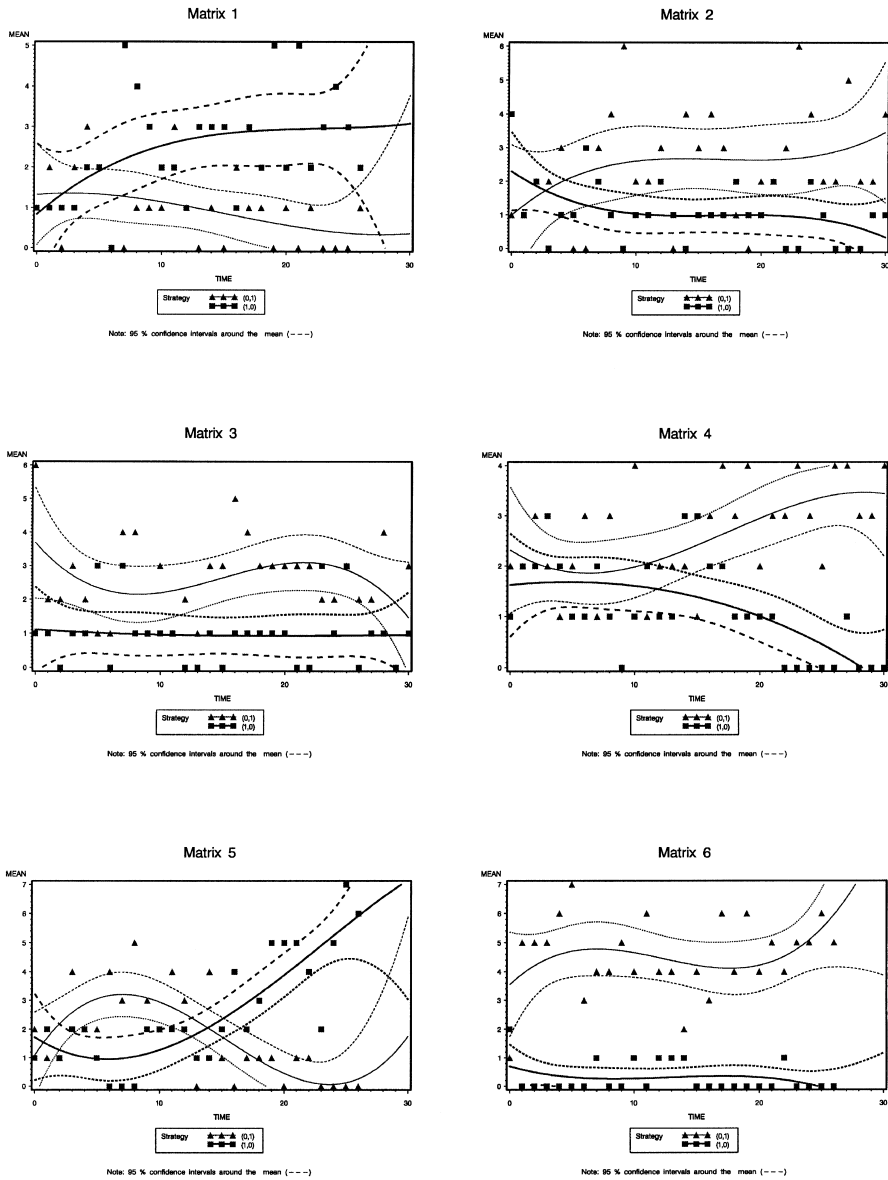


Fig. 2. Frequency of equilibria: original experiments.

discusses the evidence for learning in more detail, but we will explain here why it is an important force in favor of risk dominance.

Suppose the game is played repeatedly, as in the experiment, and over time

Table 4

Sign rank statistic: H_0 is that the difference between the number of equilibria (0,1) minus the number of equilibria (1,0) is greater than 0

Periods	Number of observations	<i>P</i> value
Trial	24	0.15571
1–5	60	0.00017
6–10	60	0.00000
11–15	60	0.00029
16–20	60	0.00056
21–25	58	0.00056
26–30	29	0.00000

players play more often the strategies that are a best response to some average of past behavior. Then both equilibria of our game are asymptotically stable, that is, the behavior converges to them if initial conditions are sufficiently close. But the risk dominant equilibrium is, by definition, the one with the largest best-response area. Because of this, it attracts the dynamics for a larger set of initial conditions. If we think of the initial conditions as random draws from some distribution, and this distribution is not biased in favor of any equilibrium, play will converge more often to the risk dominant one. Notice that if the distribution of initial conditions is constant as we increase the degree of risk dominance, learning also explains why the degree or risk dominance helps to predict the variation of behavior across games. The reason is that the degree of risk dominance measures the relative sizes of the best response areas of the different equilibria, and therefore the relative sizes of the sets of initial conditions that favor the different equilibria.

Even though adaptive behavior can go a long way towards explaining the results of this experiment, there is some unexplained coordination of initial beliefs. Table 3 shows that aggregating the experiments there is a significantly higher proportion of experiments where the initial outcome is predominantly the risk dominant one. There is a 35.6% of pairs of players with a (0,1) outcome, 17.2% with (1,0), 23.0% with (0,0) and 24.1% with (1,1). As one would expect there is a lot of initial miscoordination (actually the miscoordination persists, and we provide some reasons for this below) but clearly more people play as if they expected the risk-dominant outcome. So it is possible that some subjects were indeed trying to behave as the Bayesian rational players of Harsanyi and Selten (1988). Of course, one should take into account that their model does not predict the behavior of Bayes rational players when some other players are not Bayes rational. And, as we will show in Section 6, the model of fictitious play, which is close in spirit to Bayesian rationality, performs worse than a more naive model of learning by reinforcement.

Another possible explanation for initial coordination is that the subjects are trying to maximize the sum of payoffs at equilibrium. This would actually make a

Table 5
 Test for a linear trend increasing with the degree of asymmetry

Periods	Mean Difference						Raw <i>P</i> value	Adjusted <i>P</i> value
	Matrix 1	Matrix 2	Matrix 3	Matrix 4	Matrix 5	Matrix 6		
Trial	0.50	−1.00	2.25	0.00	0.67	3.00	0.0914	0.2000
1–5	0.20	−1.27	2.40	0.40	1.53	1.25	0.0001	0.0001
6–10	−2.00	0.60	1.70	1.60	2.53	4.90	0.0001	0.0001
11–15	−1.40	−0.73	2.40	1.20	2.27	4.00	0.0001	0.0001
16–20	−2.80	−0.80	2.40	4.00	1.00	5.40	0.0001	0.0001
21–25	−2.40	0.07	1.13	4.00	1.00	4.80	0.0001	0.0001
26–30	0.00	2.56	1.40	0.00	3.60	6.20	0.0001	0.0001

lot of sense in an experiment where subjects roles could change after each period, since that would select not only the equilibrium with the larger expected payoff for the pair, but also for each individual, at each period of time before the roles are assigned. But in our design, the roles were maintained constant precisely to reduce the role of that coordination principle. It is nevertheless possible that some subjects operate according to some kind of social convention that tells them to maximize total payoffs if no other principle can produce an acceptable solution for the game.

Some subjects may also be trying to force the occurrence of their favorite equilibrium. Table 3 shows that while type *L* agents (those endowed with an advantage) choose strategy 0 (high quality) almost all the time (from matrix 3 upwards never less than 70% of the time) type *B* agents choose action 0 more than 40% of the time except for matrix 6. This also explains why there is so much miscoordination (especially at the outcome (0,0)) in this game. We observe that except for matrix 6, equilibrium behavior is not observed in more than 60% of aggregate play, which although significantly different from the 50% expected play with random choice is still rather low coordination. This behavior is consistent with agents of type *B* that are ‘punishing’ or trying to teach the leaders to play their favorite equilibrium and is reminiscent of similar results in bargaining experiments (see, Gütz et al., 1994, Binmore et al., 1989 and Roth’s survey in Kagel and Roth, 1995) which are distantly connected with this game.

Another obvious explanation is that they may be bringing into the experiment the experience obtained with similar games played outside the laboratory. This is a hard thing to test, but we obtain limited evidence in favor of this hypothesis by having agents play two games successively with two different matrices.¹⁷ This means that experience with a different matrix is an important factor in explaining behavior but it is not determinant; the degree of risk dominance is important even with experienced subjects.

6. Learning models

Although there is some evidence that subjects’ expectations were in part coordinated on the risk dominant equilibrium at the beginning of play, Fig. 2 shows that there is a definite tendency for the equilibrium which is initially played to be reinforced over time. There are many models of learning that can explain why the risk-dominant equilibrium occurs more often. Kandori et al. (1993) show that a learning process that tends to favor strategies that are best-responses to past play, plus random mistakes implies that the risk-dominant equilibrium will be observed a larger proportion of the time. This result is quite robust, as a number of other papers, with alternative assumptions about the learning process (discrete

¹⁷For the details of the results of this test, see footnote 4.

versus continuous time, random mistakes versus aggregate shocks, long versus short memory) have also obtained it (see, e.g. Canning, 1992, Fudenberg and Harris, 1992, Di Gioacchino, 1992, Li, 1994 and Sefton, 1998).¹⁸

Since many models of learning give the same qualitative equilibrium selection conclusions in this game, finding a learning model that fits the data well is not such an important issue as in the Ultimatum game, for example (see on this, Gale et al., 1995 and Roth and Erev, 1995). It is nevertheless important to establish whether learning can, in fact, explain the characteristic features of the data better than an alternative model without learning.

A reasonable benchmark rule with no learning could be to assume that agents always play the same mixed strategy. In that case, the strategy can be estimated with the average observed play of an agent. We then estimate some learning models. We compare how well the different models fit the data using the Quadratic Deviation Measure (QDM), which is basically a sum of squared errors. Selten (1998) has recently provided an axiomatic characterization for this measure of goodness of fit. He shows that its expected value is minimized by the correct theory, but that it is less sensitive to outliers than, for example, maximum likelihood. Tang (1996) and Chen and Tang (1998) use this methodology and learning models for other games, and their results are similar to ours in terms of which models produce the smaller quadratic deviations.

Let $p_{ij}(t)$ be the probability with which an agent i of type j chooses strategy 1 at time t . Let $x_{ij}(t)$ be the actual choice of agent i of type j at time t .

The quadratic deviation of agent i of type j at time t is:

$$QDM_{ij}(t) = (x_{ij}(t) - p_{ij}(t))^2$$

Let I be the number of players and T be the number of periods in a given session. The average quadratic deviation for the players is:

$$QDM = \sum_{i=1}^I \sum_{j=A,B} \sum_{t=1}^T \frac{(x_{ij}(t) - p_{ij}(t))^2}{2I}$$

The benchmark model (B) is the *individual stationary mixed-strategy* model. This benchmark predicts that agents will choose a constant randomization which can be estimated as the actual observed frequency of play, that is, we propose that $p_{ij}(t) = \bar{p}_{ij}$ where the estimate of \bar{p}_{ij} is $\hat{p}_{ij} = \sum_t x_{ij}(t)/T$.

The first learning model we have fitted to the data is a *learning by reinforcement* (LR) model (see Roth and Erev, 1995 and Börgers and Sarin, 1997). Define by $u_{ij}(t)$ the payoff actually received by agent i of type j at period t , $x_{ij}^k(t)$ is an

¹⁸There are, however, a number of papers that reach different conclusions when there is a conflict between Pareto dominance and risk dominance (see, e.g. Binmore et al., 1995 or Robson and Vega-Redondo, 1996), which is not our case, or if the random mistakes are biased systematically in favor of one of the equilibria, as in Bergin and Lipman (1996).

indicator function that is 1 if strategy k was used by agent i of type j at time t and 0 otherwise and $R_{ij}^k(t)$ is the ‘propensity’ to play strategy k :

$$R_{ij}^k(t) = qR_{ij}^k(t - 1) + x_{ij}^k(t)u_{ij}(t)$$

where q is a parameter that measures how past information is discounted.

Notice that a strategy which is not played does not increase its ‘propensity’. The learning model predicts that strategy 0 will be played with probability:

$$p_{ij}(t) = \frac{R_{ij}^0(t)}{R_{ij}^0(t) + R_{ij}^1(t)}$$

We estimate the parameter q , as well as the initial propensities $R_{ij}^k(0)$, with a grid from 0 to 1. The higher the initial propensities, the more important initial play in the resulting outcomes. Table 6 shows that the values of the q parameter are intermediate, showing that there is some ‘forgetting’ going on but it is not complete. Consistently with this information, the initial propensities are not very high.

The other model we estimate is *modified fictitious play with experimentation* (MFP). Let $x_{ij}^0(t)$ be an indicator function that is 1 if strategy 0 was used by the opponent of agent i of type j at time t and 0 otherwise. We assume that each agent i of type j forms beliefs about the probability that her opponent will use strategy 0, which we denote by $q_{ij}(t)$:

$$q_{ij}(t) = (1 - \lambda(t)) q_{ij}(t - 1) + \lambda(t) x_{ij}^0(s)$$

where $\lambda(t) = \alpha + \beta/t$. Let $BR(q_{ij}(t))$ be the set of best pure responses to $q_{ij}(t)$. Then:

Table 6
Results from the learning models

Session number	Matrix	QDM(B)	QDM(LR)	$q(LR)$	QDM(MFP)	$\alpha(MFP)$	$\beta(MFP)$	$\delta(MFP)$
1	5	12.9	10.5	0.60	11.1	0.4	0.4	0.34
2	2	12.6	6.4	0.45	8.4	1	0	0.22
3	6	13.6	8.3	0.65	9.7	1	0	0.26
4	3	11.5	8.4	0.45	11.3	0.7	0.2	0.44
5	1	15.6	11.2	0.70	12.3	0.1	0.7	0.44
6	4	13.2	10.1	0.65	11.7	0.9	0	0.38
7	5	18.2	8.5	0.55	12.0	0.8	0	0.28
8	2	16.9	11.3	0.60	14.2	0.7	0.2	0.44
9	6	17.5	6.4	0.55	9.9	0.9	0	0.20
10	5	17.6	10.8	0.70	14.5	0.8	0.2	0.42
11	2	17.4	14.7	0.75	14.9	0.7	0	0.46
12	3	18.0	13.2	0.80	14.9	0.9	0	0.46
13	3	16.2	11.0	0.60	14.7	0.7	0.3	0.44

$$p_{ij}(t) = \begin{cases} 1 - \delta & \text{if } BR(q_{ij}(t)) = 0 \\ \delta & \text{if } BR(q_{ij}(t)) = 1 \\ \gamma & \text{if } BR(q_{ij}(t)) = \{0,1\} \end{cases}$$

so that we have fictitious play strictly speaking when $\alpha = 0$, $\beta = 1$, and $\delta = 0$; and best response dynamics when $\alpha = 1$, $\beta = 0$, and $\delta = 0$. The parameter δ adds a measure of randomness which can accommodate evolutionary stories as in Kandori et al. (1993) and Young (1993).

We estimate the parameters α, β and δ as well as the initial beliefs $q_{ij}(0)$, with a grid from 0 to 1.¹⁹ As before, we find some, but not complete ‘forgetting’, which means that neither fictitious play nor best response dynamics are totally adequate, but rather something intermediate. Also, there is a substantial amount of experimentation, which might also be a signal of individual heterogeneity.

As can be easily seen from Table 6, the benchmark model performs worse (has a higher QDM as shown by the column labeled QDM(B)) than either of the learning models that we estimate (QDM(LR) and QDM(MFP) for the learning by reinforcement and modified fictitious play models respectively). This is so, in spite of the fact that the benchmark has more parameters (one per individual) than the others (it would not be very surprising if the benchmark model performed worse than the others, if besides being more rigid it had less parameters). The modified fictitious play has a quite higher QDM than the reinforcement model, which has fewer parameters, for all sessions. These results reproduce quite closely the ones of Chen and Tang (1998) and Tang (1996), and warrant the conclusion that fictitious play is not a good learning model in this game and that the reinforcement model seems more adequate.²⁰

It is legitimate to ask how the finding that learning explains a big part of the success of risk dominance affects the motivation of the paper in terms of a model of vertical product differentiation. Do firms in real life play repeatedly the same game against opponents with whom they are randomly matched? Do they adjust play in the naive form that works well to model behavior in the lab? The answer to the first question is probably no, but one has to bear in mind that firms do introduce products in different markets at different times against different players, and we have observed in our experiment that players use their experience from one game when playing a similar one. The second question is harder to answer within the context of these experiments, but even if real life players used more sophisticated learning rules, the outcome would probably be not too different.

¹⁹The players were never indifferent between actions 0 or 1, so γ did not have to be estimated.

²⁰Sefton (1998) estimates another variant of fictitious play with mistakes and he obtains adequate results to explain the experimental data of Clark et al. (1999). He does not compare his results with those for other learning models, however.

Indeed, we have seen that the results are not very sensitive to the particular learning rule.

It is difficult to see precisely what learning rules are used by players in real life. Such an analysis was clearly beyond the scope of the present paper, where we have contented ourselves with asking whether learning might offer a motivation for behavior consistent with the prediction of the risk dominance criterion.

7. Conclusions

This paper describes the experimental evidence we have gathered on the power of risk dominance to select among the equilibria of a coordination game. This game is derived from and motivated by a vertical product differentiation model with two firms. In the theoretical model risk dominance selects the equilibrium where the high quality is produced by the firm which has some advantage.

We find that the risk dominance criterion is supported by the behavior of the experimental subjects, in the sense that the strategies that form part of such an equilibrium are selected more often than the others. Furthermore, the predictive power of the risk dominance criterion increases with the size of the initial asymmetries.

These results suggest that further research should be devoted to risk dominance, as a tool to increase the predictive power of theory, even if (as in our case) it is not a perfect predictor of the outcome. An important question is whether the result that the risk dominant equilibria happen more often than other equilibria is due to prior introspective reasoning by rational agents as Harsanyi and Selten (1988) originally proposed or whether it is the result of evolution and learning, as in Kandori et al. (1993) or Young (1993). We have partly addressed this issue through our analysis of learning in this game, but there is an alternative. In the game we have studied the two approaches (learning and introspection) give the same answers, but if there are games where the two do not lead to the same predictions, more experiments could help to choose between them.

More research on the general topic of equilibrium selection for games like the ones we study is also necessary. Harsanyi (1995) proposes a new theory of equilibrium selection which is connected to risk dominance, but permits multilateral comparisons. Selten (1995) studies a new definition of risk dominance for games with two strict equilibria where the difference in payoffs of strategies of one player are linear in the mixed strategies of other players. He shows that the measure is (essentially) the only one that satisfies a set of axioms and applies the criterion for some games. Both theories will lead to the same predictions in 2×2 games like the ones on which we did experiments, but they lead to different predictions in general. Again, more empirical evidence is needed to discriminate between the concepts.

Acknowledgements

The financial support of the Spanish Ministry of Education through DGICYT grants PB96-0302 (Cabrales), PB95-0980 (García-Fontes) PB94-0663 (Motta) and Generalitat de Catalunya grant 1997SGR00138 is gratefully acknowledged. We received helpful comments from Klaus Abbink, Fabrizio Germano, Greg Pollock, Walter Zucchini, the editor of this issue and two anonymous referees.

Appendix A. Instructions for the game (translated from Spanish)

General

You are going to participate in an experiment on an economic game. Just for participating you will receive 500 pesetas.²¹ Besides that, if you follow the instructions carefully and you take adequate decisions you may win a considerable amount of money which will be paid to you in cash after the experiment. It is important that you do not communicate with your neighbors. If you have questions raise your hand and one of us will come to you and will answer your questions.

The experiment will consist of a series of decision periods. In each period you will be randomly and anonymously matched with another person and depending on the action you choose and the action chosen by the person with whom you are matched you will receive some points. After receiving those points they will be converted into an amount of pesetas which will depend on the number of points you have. We will begin by explaining the conversion of points into pesetas so that you understand the effect of the points obtained in the amount of pesetas earned. After that we will explain the choice of actions in detail so that you understand how to obtain the points.

*Conversion of points into pesetas*²²

At the end of each period you will have obtained an amount of points between 0 and 1000 according to the system we will explain later. The computer will randomly choose an integer number between 1 and 1000 (giving all numbers the same probability). If the number of points you have is larger than or equal to the number obtained by the computer, you win 50 pesetas. If the number of points you have is smaller than the number obtained by the computer, you win 0 pesetas. For example, if you have 450 points, you have a 45% probability of winning 50

²¹ Spanish currency.

²² This paragraph obviously did not appear in the sessions where we did not correct for risk aversion.

pesetas. It is important that you realize that the larger the number of points you obtain, the larger the probability of winning.

Choice of actions

There are two types of players. At the beginning of the experiment the computer will tell you which type of player you are. You will keep this type for the duration of the experiment.

The experiment will be composed of 30 decision periods. In each period you will be randomly matched with a player of the other type. In this way all pairs will be composed of a type A player and a type B player. The identity of your couple is unknown to you and it changes every period randomly.

Once the period begins each player can choose between two options: 0 and 1. The points obtained will be a function of your choice and the independent choice of your couple, as indicated in the following example:

		Choice of B	
		0	1
Choice of A	0	292, 164	612, 362
	1	364, 509	367, 346

For instance, to know the points you have obtained if you are type A and you choose 0 when the person you are matched with (therefore a type B) chooses 1, you look in the table for the cell corresponding to (A, 0), (B, 1) and you find the number 612. In this same situation your pair obtains 362, as one can see from the same cell (A, 0), (B, 1) of the table. You can find the table for this experiment at the end of these instructions.

Once you decide on your action you should enter your decision in the computer.

Once you and the player you are matched with have taken the decision, the computer will determine the number of points earned, using the corresponding table. The computer will tell you this information.

After that, the computer converts the points into pesetas. Once this is done the computer will tell you this information together with the history of your decisions and those of the player with whom you have been matched.

History of decisions

The computer gives you a history of the decisions you have taken in previous periods, as well as those of the players you were matched with and how many points and pesetas you won.

Development of the experiment

At the beginning of the experiment you will have two *practice* rounds to get used to the terminal and the game. The gains of those periods are fictitious and will not be paid at the end. Once the *practice* rounds are over the true experiment begins, where the gains are real. The experiment will consist of 30 *true* periods.

Appendix B. Details for the statistical tests

For the sign rank test reported in Table 4, we first computed the total number of (0,1) and (1,0) equilibria and we then computed the difference between these two scores for each experiment and each period. We then grouped the periods in sequences of 5, and computed the sign rank statistic as $S = p - n/2$, where p is the number of values greater than 0 and n is the number of nonzero values. Under the null hypothesis that the median of the difference is zero, the probability that the sign statistic is greater or equal than the observed value is $2 \sum_{j=0}^{\min\{p,n-p\}} \binom{n}{j} 0.5^n$.

For the linear trend test reported in Table 4, we computed the difference as in the previous test, and we computed the centered trend coefficients, which are then used to form a t -statistic to contrast the within-group means. The P values were adjusted using a bootstrap method.

References

- Bergin, J., Lipman, B., 1996. Evolution with state-dependent mutations. *Econometrica* 64, 943–956.
- Binmore, K., 1992. *Fun and Games*, D.C. Heath, Lexington, MA.
- Binmore, K., Samuelson, L., Vaughn, R., 1995. Musical chairs: Modelling noisy evolution. *Games and Economic Behavior* 5, 387–424.
- Binmore, K., Shaked, A., Sutton, J., 1989. An outside option experiment. *Quarterly Journal of Economics* 104, 753–770.
- Börgers, T., Sarin, R., 1997. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 77, 1–14.
- Brezis, E., Krugman, P.R., Tsiddon, D., 1993. Leapfrogging in international competition: A theory of cycles in national technological leadership. *American Economic Review* 83, 1211–1219.
- Budd, C., Harris, C., Vickers, J., 1993. A model of the evolution of duopoly: does the asymmetry between firms tend to increase or decrease? *Review of Economic Studies* 60, 543–573.
- Cabrales, A., Motta, M., 1996. Country asymmetries, endogenous product choice and the speed of trade liberalization, CEPR, London, WP 1326.
- Canning, D., 1992. Average behavior in learning models. *Journal of Economic Theory* 57, 442–472.
- Carlsson, H., Van Damme, E., 1993. Global games and equilibrium selection. *Econometrica* 61, 989–1018.
- Charness, G., 1998. Pre-play communication and credibility: A test of Aumann's conjecture, Universitat Pompeu Fabra, Barcelona, Working Paper 293.
- Chen, Y., Tang, F.F., 1998. Learning and incentive compatible mechanisms for public goods provision: An experimental study. *Journal of Political Economy* 106, 633–662.

- Clark, K., Kay, S., Sefton, M., 1999. When are Nash equilibria self-enforcing? An experimental analysis. *International Journal of Game Theory* (forthcoming).
- Cooper, R., De Jong, D.V., Forsythe, R., Ross, T.W., 1989. Communication in the battle of sexes games: Some experimental results. *RAND Journal of Economics* 20, 557–568.
- Cooper, R., De Jong, D.V., Forsythe, R., Ross, T.W., 1993. Forward induction in the battle of sexes game. *American Economic Review* 83, 1303–1316.
- Dasgupta, P., Stiglitz, J., 1988. Learning-by-doing, market structure and industrial and trade policy. *Oxford Economic Papers* 40, 246–268.
- Di Gioacchino, D., 1992. Learning and equilibrium selection in 2×2 games. *Economics Letters* 38, 405–409.
- Farrell, J., Shapiro, C., 1988. Dynamic competition with switching costs. *RAND Journal of Economics* 19, 123–137.
- Flam, H., Helpman, E., 1987. Vertical product differentiation and North–South trade. *American Economic Review* 77, 810–822.
- Friedman, D., 1996. Equilibrium in evolutionary games: Some experimental results. *Economic Journal* 106, 1–25.
- Fudenberg, D., Harris, C., 1992. Evolutionary dynamics in games with aggregate shocks. *Journal of Economic Theory* 57, 420–442.
- Gale, J., Binmore, K., Samuelson, L., 1995. Learning to be imperfect: The ultimatum game. *Games and Economic Behavior* 8, 56–90.
- Grossman, G., Helpman, E., 1991. Quality ladders and product cycles. *Quarterly Journal of Economics* 106, 557–586.
- Gütz, W., Schmittberger, R., Schwarze, B., 1994. An experimental analysis of ultimatum bargaining. *Journal of Economic Behaviour and Organization* 3, 367–388.
- Guyer, M., Rapoport, A., 1972. 2×2 games played once. *Journal of Conflict Resolution* 62, 409–431.
- Harsanyi, J., 1995. A new theory of equilibrium selection for games with complete information. *Games and Economic Behavior* 8, 91–122.
- Harsanyi, J., Selten, R., 1988. *A General Theory of Equilibrium Selection in Games*, The MIT Press, Cambridge, MA.
- Kagel, J., Roth, A.E., 1995. *Handbook of Experimental Economics*, Cambridge University Press, New York.
- Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation and long-run equilibrium in games. *Econometrica* 61 (1993), 29–56.
- Li, S., 1994. Dynamic stability and learning processes in 2×2 coordination games. *Economics Letters* 46, 105–111.
- Luce, R., Raiffa, H., 1957. *Games and Decisions*, Wiley, New York.
- Motta, M., Thisse, J.F., Cabrales, A., 1997. On the persistence of leadership and leapfrogging in international trade. *International Economic Review* 38, 809–824.
- Mussa, M., Rosen, S., 1978. Monopoly and product quality. *Journal of Economic Theory* 18, 301–317.
- Reinganum, J., 1985. Innovation and industry evolution. *Quarterly Journal of Economics* 100, 81–99.
- Robson, A., Vega-Redondo, F., 1996. Efficient equilibrium selection in evolutionary games with random matching. *Journal of Economic Theory* 70, 65–92.
- Roth, A., Erev, I., 1995. Learning in extensive games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8, 164–212.
- Roth, A., Malouf, M.W.K., 1979. Game-theoretic models and the role of bargaining. *Psychological Review* 86, 574–594.
- Sefton, M., 1998. A model of behavior in coordination game experiments, University of Manchester, Technical report (mimeo).
- Selten, R., 1995. An axiomatic theory of a risk dominance measure for bipolar games with linear incentives. *Games and Economic Behavior* 8, 213–263.
- Selten, R., 1998. Axiomatic characterization of the quadratic scoring rule. *Experimental Economics* 1, 43–62.

- Selten, R., Sadrieh, A., Abbink, K., 1995. Money does not induce risk neutral behavior, but binary lotteries do even worse, Universität Bonn, Discussion Paper B-343 (SFB 303).
- Shaked, A., Sutton, J., 1982. Relaxing price competition through product differentiation. *Review of Economic Studies* 49, 3–13.
- Straub, P., 1995. Risk dominance and coordination failure in static games. *Quarterly Review in Economics and Finance* 55, 339–363.
- Tang, F., 1996. Anticipatory learning in two-person games: An experimental study. Part II. Learning, Universität Bonn, Discussion Paper B-362 (SFB 303).
- Van Huyck, J., Battalio, R., Beil, R., 1990. Tacit coordination games, strategic uncertainty and coordination failure. *American Economic Review* 80, 2234–2249.
- Young, P., 1993. The evolution of conventions. *Econometrica* 61, 57–84.