Risk dominance selects the leader. An experimental analysis *

Antonio Cabrales †
Walter García-Fontes †
Massimo Motta †

February 1997

*The financial support of the Spanish Ministry of Education through DGYCIT grants PB93-0398 (Cabrales), DI000526 (García-Fontes) and PB94-0663 (Motta) are also gratefully acknowledged. We received helpful comments from Klaus Abbink, Fabrizio Germano, Greg Pollock.

†Department of Economics, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona (Spain)

Abstract

Coordination games arise very often in studies of industrial organization and international trade. This type of games has multiple strict equilibria, and therefore the identification of testable predictions is very difficult. We study a vertical product differentiation model with two asymmetric players choosing first qualities and then prices. This game has two equilibria for some parameter values. However, we apply the risk dominance criterion suggested by Harsanyi and Selten and show that it always selects the equilibrium where the leader is the firm having some initial advantage. We then perform an experimental analysis to test whether the risk dominance prediction is supported by the behaviour of laboratory agents. We show that the probability that the risk dominance prediction is right depends crucially on the degree of asymmetry of the game. The stronger the asymmetries the higher the predictive power of the risk dominance criterion.

1 Introduction

In this paper we analyze a vertical product differentiation model where two firms first decide on the quality they want to bring to the market and then choose prices. This 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 at 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. The first where the leader is the firm enjoying the initial advantage, and the second where the leader the firm starting with an initial disadvantage.

Several interpretations for this model and for the initial asymmetries are possible. For instance, firms might differ in the technology available to them, so that the initial advantage would consist of a lower cost of carrying out the research and development activities necessary to improve the quality of the good. Another interpretation 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 update the quality of their product up to the desired level. Under the latter 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.

Whatever the nature and interpretation which might be given to the asymmetries, one could wonder whether the result that when initial asymmetries are not too large the leader cannot be predicted survives the application of some equilibrium selection criteria. Standard criteria of selection among multiple equilibria (such as perfectness, properness, strategic stability and Pareto dominance) have no selective power in the game we study. For this reason we resort to the concept of risk dominance proposed by Harsanyi and Selten (1988). We apply this criterion and we show that 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. Nevertheless, when the firms are identical risk dominance does not select a leader. The presence of such discontinuity (unambiguous selection for small asymmetries but no selection for symmetric firms) suggests that we should be cautious in applying the result when the initial asymmetries are small. As we will see later, the experimental analysis supports this view.

To test the prediction of the risk dominance criterion we have performed experiments on a leadership game derived from our vertical product differentiation model. Since we wanted to concentrate on the equilibrium selection problem we designed a game that shares the characteristics of the theoretical model proposed but whose structure was easier to understand by the experimental subjects. We used a game in which only two actions are available to each player, thus proposing a strong discretization of the strategy space (the continuum in the original differentiation model), and the payoffs are those of the equilibrium strategies (suitably renormalized) for the differentiation model. The resulting game is a version of the Battle of the Sexes, or Chicken,[1] with different degrees of asymmetry (see Table 1; matrix 1 is symmetric, while matrices 2 to 6 show increasing asymmetries in payoffs).

The main hypotheses we tested were whether the equilibrium where the firm with initial advantage is the leader, or the equilibrium where the disadvantaged firm becomes the leader were selected by the subjects in this game, and whether the answer depends on the initial degree of asymmetry. Since the equilibrium with leadership by the firm with an advantage is risk dominant this can be viewed additionally as a test of the empirical relevance of risk dominance as an equilibrium selection criterion. 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 the evolutionary models that have suggested risk dominance as an equilibrium selection criterion are useful in explaining the facts.

The experimental results support the hypothesis that leadership by the firm with an advantage is the result one should expect more often and therefore gives support for the use of risk dominance as an equilibrium selection criterion in this type of games. But this support has to be qualified 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 also tested different learning protocols and we found that fictitious play does a worse job at explaining the data than unsophisticated stimulus-response models which turn out to be related to evolutionary processes. In addition to their intrinsic game theoretic interest, we wanted to evaluate different learning rules in this game

[1] One can find discussions of such games in any standard game theory textbook, like Binmore (1992) or Gibbons (1992)

because some recent research (Kandori, Mailath and Rob 1993, Young 1993) has shown that 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 in a coordination game like the one we are analyzing at the beginning of play (and in the absence of strong focal points), this is a bit surprising when it lasts for 25 to 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 in, 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.

Our paper is related to two different strands of the economic literature. The first is the wide literature in industrial organisation 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, Harris and Vickers (1993). Brezis, Krugman and Tsiddon (1993), Flam and Helpman (1987) and Motta, Thisse and Cabrales (1997) study the problem from an international perspective, and model situations where 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 is Guyer and Rapoport (1972). They studied a number of 2×2 games, and among them there was an asymmetric version of Chicken. They find that the actions leading to the maximin equilibrium in that game are selected more often than the others and that the frequency of these choices increases with the asymmetry of the payoffs. Their results are consistent with ours in the sense that the maximin equilibrium in that game is also the risk dominant equilibrium so the risk dominant equilibrium is chosen more often. However, the fact that in their experiment, unlike in ours, the risk dominant equilibrium is also the maximin equilibrium, makes it difficult to assess if the result is due to risk dominance or to max-

[2] See v.g. Friedman (1996).

imin. One has to remember that at the time those experiments were performed the concept of risk dominance was not known, so there was little reason to conduct the experiments with games that separated risk dominance from maximin as we do. Another closely related experimental paper is Friedman (1996). He studies convergence to equilibrium in a variety of 2×2 games, including an asymmetric version of the Battle of the Sexes. He finds that players converge to playing equilibrium profiles a large percentage of the times, but he does not report to which equilibria. One important difference with our design is that he does not vary the degree of asymmetry in the payoffs.

Other previous experimental research on similar games has tended to focus on symmetric games. Cooper, DeJong, Forsythe and Ross (1989) run experiments with the symmetric Battle of the Sexes with preplay communication to test Farrell's (1987) theory of cheap talk and equilibrium selection. They show that communication, especially the one-sided variety, is an empirically useful equilibrium selection device. Cooper, DeJong, Forsythe and Ross (1993) do experiments with the symmetric Battle of the Sexes plus an outside option (thus the complete game is not symmetric) to test forward induction as an equilibrium selection device. They show that the outside option does influence the outcome of the game but that this is not solely due to forward induction but also a consequence of the creation of a focal point through the asymmetry of the game. Van Huyck, Battalio and Beil (1990) study a symmetric coordination game with 7 strategies (levels of effort), where all the homogeneous strategy profiles are equilibria. They find that the risk dominant equilibrium (the minimum effort)[3] is eventually played by all subjects, even though it is the most inefficient of all equilibria.

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 experiment 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

Let us 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

[3] Sacco (1996) shows that the equilibrium with all agents playing the minimum effort in this game is risk dominant.

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. We assume the distribution of θ to be uniform with $\theta \in [0, \bar{\theta}]$ and a density S.

We assume that there exist two firms in the industry, A and B. In the first period of the game they decide on the quality they want to produce u_j and incur a fixed (i.e. independent of the quantity produced) cost of quality $F_j = k_j \frac{u_j^2}{2}$, with $j = A, B$. The cost of quality can be thought of as R&D investments or advertising outlays. The quadratic form we have taken is a standard assumption which greatly simplifies the calculations. Any convex function would give the same qualitative results. Note that the two firms do not necessarily have the same technology: firm A is at least as efficient as firm B, with $1 = k_A \leq k_B = k$. 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 decide on the price at which they want to sell their product. We work as usual by backward induction, and solve the last stage of the game first.

We first find the demand schedules faced by the top and bottom quality firm respectively as:

$$q_1 = S\left(\bar{\theta} - \frac{p_1 - p_2}{u_1 - u_2}\right), \qquad q_2 = S\left(\frac{p_1 - p_2}{u_1 - u_2} - \frac{p_2}{u_2}\right),$$

where $u_1 \geq u_2$. It is then straightforward to derive the first order conditions, compute the price equilibrium and check that profits at the price equilibrium for the top and low quality are:

$$\Pi_1 = \frac{4u_1^2(u_1 - u_2)S\bar{\theta}^2}{(4u_1 - u_2)^2}, \qquad \Pi_2 = \frac{u_1 u_2(u_1 - u_2)S\bar{\theta}^2}{(4u_1 - u_2)^2}.$$

At the first stage of the game, the net profit functions for the firms are given by $\pi_1 = \Pi_1 - k_j \frac{u_1^2}{2}$, and $\pi_2 = \Pi_2 - k_i \frac{u_2^2}{2}$, with $j, i = A, B$ and $j \neq i$.

Note that we have deliberately not specified which firm is producing the top and which the bottom quality, since either firm can be the high (low) quality provider at equilibrium. Indeed, there might exist two equilibria in pure strategies (we do not consider mixed strategies here). In the first one, it is the more efficient firm A which produces the top quality. In the second, it is the less efficient firm B. We now turn to the characterization of both equilibria, and to the analysis of their existence.

2.1 The more efficient firm produces the high quality

If the top quality firm is the one with lower costs of quality production (firm A), the first-order conditions of the problem are:

$$u_1 = 4S\bar{\theta}^2 \frac{4u_1^2 - 3u_1u_2 + 2u_2^2}{4(4u_1 - u_2)^3} \equiv \phi(u_1, u_2), \tag{1}$$

$$ku_2 = S\bar{\theta}^2 u_1^2 \frac{4u_1 - 7u_2}{4(4u_1 - u_2)^3} \equiv \mu(u_1, u_2). \tag{2}$$

By dividing the two equations above, rearranging and writing $u_1 = ru_2$ with $r \geq 1$ we obtain:

$$u_2{}^3(4r^3 - 16kr^2 - 7r^2 + 12kr - 8k) = 0. \tag{3}$$

The only meaningful root for this equation is:

$$r_A = \frac{7 + 16k}{12} + \frac{49 + 80k + 256k^2}{12(g(k) + 24\sqrt{3kh(k)})^{1/3}} + \frac{(g(k) + 24\sqrt{3kh(k)})^{1/3}}{12},$$

where: $g(k) = 343 + 2568k + 1920k^2 + 4096k^3$; $h(k) = 686 + 2967k + 3552k^2 + 5888k^3$.

Note that $\frac{\partial \pi_2}{\partial u_2} = 0$ can be written as: $u_2 = S\bar{\theta}^2 \frac{r^2(4r-7)}{(4r-1)^3}$. By replacing r with r_A, we obtain u_2^* and $u_1^* = r_A u_2^*$ as functions of k only (the term $S\bar{\theta}^2$ has just a scale effect throughout).

Figure 1 (left-hand panels) depicts the qualities and profits for this candidate equilibrium $E_1(u_1^*, u_2^*)$, where the top quality is provided by the more efficient firm A. Note that the top quality makes considerably higher profits than the low quality firm.

2.2 The less efficient firm produces the top quality

If firm B were the top quality firm, the first-order conditions of the problem would be:

$$ku_1 = \phi(u_1, u_2), \tag{4}$$

$$u_2 = \mu(u_1, u_2). \quad (5)$$

Let us write $u_1 = zu_2$ (with $z \geq 1$) and use the same procedure followed in the previous section to derive the solution. We can then find the value z_B which satisfies the first-order conditions, and by substitution the solution $E_2 = (u_1^{**}, u_2^{**})$. For completeness, we report here the value z_B which is:

$$z_B = \frac{7k+16}{12k} + \frac{256+80k+49k^2}{12k(l(k)+24k\sqrt{3m(k)})^{1/3}} + \frac{(l(k)+24k\sqrt{3m(k)})^{1/3}}{12k},$$

where: $l(k) = 4096 + 1920k + 2568k^2 + 343k^3$; $m(k) = 5888 + 3552k + 2967k^2 + 686k^3$.

Figure 1 (right-hand panels) reports qualities and profits at this candidate solution.

Note that the two pairs of candidate solutions have been obtained under the hypothesis that no firm can deviate from the quality it has been assigned. For instance, in the first case the candidate solution E_1 was found under the hypothesis that firm A produces the top quality, and firm B the bottom quality. But to make sure that the pair (u_1^*, u_2^*) is really an equilibrium, we also have to check that firm B does not find it profitable to 'leapfrog' the rival and provide a quality higher than (u_1^*). In other words, it must be checked that there exists no quality u_1' such that $\pi_1'(u_1', u_2 = u_1^*) \geq \pi_2^*(u_1^*, u_2^*)$. Likewise, it must be checked that firm A does not have an incentive to deviate by supplying a quality which is lower than u_2^*. Indeed, it is possible to show that these deviations are not profitable, and therefore conclude that the pair (u_1^*, u_2^*) is always an equilibrium (see Motta, Thisse and Cabrales, 1997, for an illustration in a similar model). The discussion below should give more insight about this result.

The same exercise must be made for the second case, where firm B produces the top quality. However, it turns out that this is not an equilibrium for all the values of the parameters. Indeed, there exist high enough values of parameter k (to be precise, $k = 1.5894$ is approximately the threshold value above which this equilibrium collapses) for which the more efficient firm finds it profitable to produce a quality u_1' higher than the quality u_1^{**} the rival would produce at the candidate solution. In other words, $\pi_1'(u_1', u_2 = u_1^{**})$ can be higher than $\pi_2^{**}(u_1^{**}, u_2^{**})$, as can be seen from figure 2.

To understand why this equilibrium breaks down when technological asymmetries are large enough, consider the extreme case where firm B is infinitely inefficient.

If k tends to infinity, then firm B will choose a top quality $u_1^{**} = \epsilon$, with ϵ arbitrarily small, since a huge investment must be made even to produce even a very low quality. At the candidate equilibrium, firm B is making infinitesimally small profits, and firm A's profits are even lower (the bottom quality firm always makes less profit). It is then clear that the latter firm has an incentive to deviate from the candidate equilibrium. At a small cost, it can produce a quality higher than ϵ, become the top firm and earn higher profits.

2.3 Equivalent models

We have seen above an example of a model which gives rise to two possible equilibria with asymmetric payoffs, and of which only one is risk-dominant across all the values of the parameters. However, there exist many other possible examples of vertical product differentiation models which share the same basic features.

Consider for instance the following variation of the model presented above. 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 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, and it is always risk-dominated by the former equilibrium. This model has been analyzed, in an international trade context by Motta, Thisse and Cabrales (1997)[4]

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. Standard 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 no equilibrium Pareto dominates the other (tak-

[4]See also Cabrales and Motta (1996) for another model with very similar features.

ing into account only the welfare of the players, the firms; and not the consumers). There is a solution concept that selects between equilibria in our game, though, and this is the concept of risk-dominance introduced by Harsanyi and Selten (1988).

In the first subsection we will define the concept of risk-dominance for 2×2 games, which we will use in the experiments, and 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. In the second subsection we define risk dominance for games with more strategies and we show that risk dominance also selects the equilibrium where the firm with the lower cost of quality is the leader in a game with a less coarse discretization of the strategy space of the theoretical game in the previous section.

3.1 Risk dominance in the 2×2 game

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 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 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 with the following payoff matrix.

	B_0	B_1
A_0	a_{00}, b_{00}	a_{01}, b_{01}
A_1	a_{10}, b_{10}	a_{11}, b_{11}

where the payoffs are such that $E_0 = (A_0, B_1)$ and $E_1 = (A_1, B_0)$ are strict Nash equilibria[5], and let $LA_0 = a_{01} - a_{11}$. LA_0 is the gain made by player A by predicting rightly that the other player will play as in E_0 (and best responding to the prediction) instead of predicting wrongly that the other player will play as in E_1 (and best responding to the prediction). Similarly, let $LB_0 = b_{01} - b_{00}$. $LA_1 = a_{10} - a_{00}$. $LB_1 = b_{10} - b_{11}$. We say that equilibrium E_0 risk dominates equilibrium E_1 when $LA_0 LB_0 > LA_1 LB_1$.

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 considered a good equilibrium selection criterion. In evolutionary models (like Kandori, Mailath and Rob 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. This happens because the likelihood of a stationary state depends on the size of the shock needed to escape from it, and the size of the shock needed to escape an equilibrium (the only stationary states) depends on the size of the area for which the equilibrium strategies are best responses. Risk dominant equilibria are precisely those with the largest best response areas.

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

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, like the papers of Guyer and Rapoport (1972) and Van Huyck, Battalio and Beil (1990) [6] we cite in the introduction, the risk dominant equilibrium is chosen by the subjects.

Let us now apply the risk dominance criterion to our model. Denote by E_0 the equilibrium where the low cost firm is the leader and by E_1 the equilibrium where the high cost firm is the leader. Recall that A is the low cost firm and B is the high cost firm.

In our case $a_{00} = \pi_{1A}(u_1^{**}, u_2^*)$, $a_{01} = \pi_{1A}(u_1^*, u_2^*)$, $a_{10} = \pi_{2A}(u_2^{**}, u_1^{**})$, $a_{11} = \pi_{1A}$ $(u_1^*\ , u_2^{**})$, $b_{00} = \pi_{2B}(u_1^{**}, u_2^*)$, $b_{01} = \pi_{2B}(u_1^*, u_2^*)$, $b_{10} = \pi_{1B}(u_2^{**}, u_1^{**})$ and $b_{11} = \pi_{2B}(u_1^*, u_2^{**})$.

In our case, LA_0 is what player A gains by forecasting rightly that the other player will play the equilibrium where A itself is the leader, instead of forecasting wrongly that the other equilibrium holds. LA_1 represents the gains for player A of forecasting rightly that the other player will play the equilibrium where B is the leader. The interpretation of LB_1 and LB_2 is analogous.

Figure 3 shows that we have $LA_0 \geq LB_1$ and $LB_0 \geq LA_1$, 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 lower 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 $\frac{LA_0 LB_0}{LA_1 LB_1 + LA_0 LB_0}$. 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

[6] For an interpretation of these results from an adaptive learning perspective, see Crawford (1995) and Broseta (1993)

dominated E_1 is, and so the likelier it is that we will observe players using the E_0 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.

3.2 Risk dominance in the game with more than two strategies

We consider a 2 player game, G, where the strategy space for player $i = A, B$ is U_i (in our case the qualities), and the payoff function is $\pi_i : U_A \times U_B \to \Re$.

Extending risk dominance to games with more than 2 strategies proceeds by first postulating a theory of preliminary expectations and then using the tracing procedure to arrive at one of the two equilibria from those expectations. Since the risk dominance criterion assumes that only one of two equilibrium strategy pairs will be played, only strategies that are somehow connected with the equilibrium strategies should be a part of the preliminary expectations. To do this more formally, let us define a game G' as a *formation* of G if the set $U'_A \times U'_B$ formed by reducing the strategy sets from the original game and maintaining the payoff function is closed with respect to best replies in G. In other words, G' is a *formation* of G if

$$\forall u_j \in U'_j \quad B_i(u_j) \subseteq U'_i, \quad i, j = A, B;\ j \neq i, \tag{6}$$

where B_i is the best response correspondence of agent i for game G. Let u^* and u^{**} be two equilibrium points. Let F be the smallest *formation*[7] such that u^* and u^{**} belong to the strategy sets in F. F will be the game used for the risk dominance comparison between u^* and u^{**}.

The preliminary expectations used by Harsanyi and Selten (1988) are given by the *bicentric priors*. For every z with $0 \leq z \leq 1$ and for $i = 1, 2$ define

$$r_i^z = a_i(zv_j + (1-z)v'_j), \quad j \neq i \tag{7}$$

where a_i is the *centroid* of the best response correspondence for i and z is some probability that i could assign to the equilibrium v. The *centroid* of a set U'_i is a mixed strategy c such that $c(u_i) = 1/|U'_i|$ if $u_i \in U'_i$ and $c(u_i) = 0$ otherwise; where $|U'_i|$ is the number of elements in U'_i. r_i^z is a best response for player i if she believes

[7] Such a set exists because the intersection of two formations is a formation.

j will play v_j with probability z and v'_j with probability $1 - z$. The *bicentric prior* of player j about player i strategy will be defined then as follows:

$$p_i(u_i) = \int_0^1 r_i^z(u_i)dz, \quad \text{for every } u_i \in U_i \tag{8}$$

Player j does not know what probability player i assigns to v_j and v'_j, so player j will (adopting the principle of insufficient reason) give the same weight to all possible z for player i.

These bicentric priors represent the initial expectations of each player about the other player's strategy. These expectations need not be consistent with the actual strategy the other player intends to play given her preliminary expectations, and thus the best responses to preliminary expectations need not be equilibria. The best responses to preliminary expectations will be used by the tracing procedure to start a path that will smoothly approach one equilibrium which will be the one selected by the risk dominance criterion.

The (linear) tracing procedure is defined as follows: for $0 \leq t \leq 1$, let the game Γ^t be the game with the same strategy sets as F, and whose payoff functions are defined by

$$\pi_i^t(u_i, u_j) = t\pi_i(u_i, u_j) + (1 - t)\pi_i(u_i, p_j), \qquad i, j = A, B, \, j \neq i \tag{9}$$

where π_i is the payoff function in F and p are the *bicentric priors*. Clearly $\Gamma^1 = F$ and Γ^0 is the game where each player is facing her preliminary expectations.

For any game Γ^t let E^t be the set of equilibrium points in Γ^t and let $X = X(F, p)$ the graph of the correspondence $t \to E^t$ for $0 \leq t \leq 1$. A continuous path L contained in X connecting point $x^0 = (0, u^0)$ and $x^1 = (1, u^1)$ (where u^t is an equilibrium of Γ^t, $t = 0, 1$), is called a *feasible path*. The *linear tracing procedure* consists in selecting an equilibrium of a game F which is the strategy part u^1 of the end-point $(1, u^1)$ of a *feasible path* L. The linear tracing procedure is *well defined* if X contains one and only one *feasible path*. Take a pair (F, p) for which the tracing procedure is *well defined* and let u^1 be the equilibrium selected. We denote then $T(F, p) = u^1$.

We say that an equilibrium v *risk dominates* v' if, given a *bicentric prior* p, and a reduced game F, $T(F, p) = v$.

We have computed numerically the preliminary expectations and then applied the tracing procedure for a game with payoff functions like the ones in section 2 and strategy spaces reduced to the two equilibrium strategies and one hundred convex combinations of them. We find that for all the parameter values for which we do the computation ($k = 1.1, 1.2, 1.3, 1.4, 1.5$ and notice that for $k > 1.5$ there is a unique equilibrium) the equilibrium with the low cost firm being the leader is one selected by risk dominance in all the cases. Table 2 reports the results of the tracing procedure. For selected values of t between 0 and 1[8] we show the equilibrium values of the auxiliary games. One can easily see that there is only one path connecting the equilibrium of the game with the preliminary expectations and the actual game played. This path is the constant path that already starts at the equilibrium values of quality for leadership by the low cost firm. For sufficiently high values of t another path arises, which finally connects with the other equilibrium. This level of t is higher as the value of k rises. This is similar to the rising value of the degree of risk dominance, $\frac{LA_0 LB_0}{LA_1 LB_1 + LA_0 LB_0}$, we defined for 2×2 games, and can be interpreted as saying that the equilibrium where the low cost firm is the leader is "more" risk dominant as the cost asymmetry increases.

4 Experimental design

In the experiments we chose to restrict the number of strategies to two. There are several reasons for this. Firstly, we have 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 coarser discretization of the strategy space. 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 behaviour 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 mumber of times the game was played and they had two initial practice sessions 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.

[8] We only report 10 values of t, but the computations were done for 100.

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 the instructions which can be found in the appendix 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-Malouf (1979)). Each player received the payoff in points which then determine 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 [9]. 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, Sadrieh and Abbink 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 to the experimental subjects. There were no major qualitative changes in behavior in those control sessions.

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[10] This was done to test for ex-

[9]Notice 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.

[10]When 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.

perience 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 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 to 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 A is playing the quality level 25.28 and B 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. The subjects, however, 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.

There are two pure strategy equilibria independently of the matrix used (i.e. independently of the value of k), and they are (0,1) and (1,0). The former equilibrium is the risk dominant one (except in the symmetric case of matrix 1).

Table 3 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 4 and Figure 4 summarize the results obtained with the pooled data from the 12 experiments [11]. In Figure 4, as in Figures 5, 6, 7 and 8, 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 $(0,1)$ minus the number of equilibria $(1,0)$ is greater than 0 [12]. We pooled the observations for groups of 5 periods. The results are reported in Table 5, 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 ratio $\frac{LA_0\,LB_0}{LA_1\,LB_1+LA_0\,LB_0}$, which is an increasing function of k as can be seen from Figure 3. It is clear that the proportion of times for which the risk dominant equilibrium is played increases with k as can be seen in Figures 5, 6, 7 and 8, somehow confirming the weaker interpretation of risk dominance we proposed. 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 6. The results strongly show that we cannot reject the hypothesis that the difference is linearly increasing.

We replicated the experiments based on Matrices 5 and 6, and the results can be seen in Figure 6. The replication of the experiment based on matrix 6 gives very similar results to the original. The replication of the experiment based on matrix 5 runs more in accordance with the prediction of risk dominance than the original. The only difference between the original experiment and the replication is that the total number of periods was extended from 25 to 30, but the behavior at period 25 is already different enough to consider improbable that this is cause for the difference. Another experiment with matrix 5 was performed without using the binary lottery procedure. The results, which can be seen in Figure 8 also show a clear prevalence of the risk dominant equilibrium.

We also checked for experience effects, and the results are shown in Figure 7. We run an experiment with matrix 2 based on subjects experienced on matrix 5. In contrast with the unexperienced subjects that are confronted with matrix 2, which

[11] Figure 4 also includes the two trial periods performed in each session

[12] The details for this and all subsequents tests are shown in the Appendix B.

tend to produce an outcome in favor of the risk dominant equilibrium (see Figure 5), in this case the outcome goes in favor of the $(1,0)$ equilibrium as in the original experiment with matrix 5. Matrix 3 is run with subjects experienced in matrix 6. In this case, both the experienced and the unexperienced subjects, chose more often the risk dominant equilibrium, but the subjects who had experienced matrix 6 did it even more often than the unexperienced players. Matrix 4 is run with subjects experienced in matrix 1. In the two cases where the matrix was run with and without experience (matrices 2 and and 3) the subjects behaved significantly differently in the experience treatment than in the original treatment, and in the experience treatment they moved towards the outcome observed in the previous game they game had played, but their behavior was not identical as in the previous game. 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.

Finally we check for the effects of the lottery mechanism that we used in the experiments. Comparing Figure 8 with the original matrix, we can see that the prevalence of the risk dominant equilibrium seems to get reinforced if the lottery mechanism is not used. Note, however, that strictly speaking what we observe without the lottery procedure is that the equilibrium which would be risk dominant under risk neutrality is played more often than with the lottery procedure.

Table 4 also shows that that while type A agents (those endowed with an advantage) choose strategy 0 (high quality) almost all the time (from matrix 3 upwards never less than 70 percent of the time) type B agents choose action 0 over 40 percent of the time except for matrix 6. This is consistent with their "punishing" or trying to teach the leaders to play their favorite equilibrium and is reminiscent of similar results in bargaining experiments (see Güth, Schmittberger and Schwarze 1982, Binmore, Shaked and Sutton 1989 and Roth's survey in Kagel and Roth 1995) which are distantly connected with this game.

6 Learning models

Since some of the theoretical models that lend support to risk dominance as an equilibrium selection criterion are learning or evolutionary models we would like to present here some evidence as to whether agents learn in this game and which model fits the data best.

As a benchmark rule, we assume that agents could play always the same mixed strategy. In that case, the strategy can be estimated with the average observed play of an agent. Then we estimate other learning models and we obtain the Quadratic Deviation Measure (QDM), which is basically a sum of squared errors. Tang (1996) and Chen and Tang (1996) 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 $p_{ij}(t)$ 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 $\bar{p}$ is $\hat{p} = \sum_t x_{ij}(t)/T$.

The first learning model we have fitted to the data is a *learning by reinforcement* (LR) model (see Roth and Erev 1994, Börgers and Sarin 1997). Define by $u_{ij}(t)$ the payoff actually received by agent i of type j at period t, $x^k_{ij}(t)$ is an indicator function that is 1 if strategy k was used by agent i of type j at time t and 0 otherwise and $R^k_{ij}(t)$ is the "propensity" to play strategy k.

$$R^k_{ij}(t) = qR^k_{ij}(t-1) + x^k_{ij}(t)u_{ij}(t)$$

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

$$p_{ij}(t) = \frac{R^0_{ij}(t)}{R^0_{ij}(t) + R^1_{ij}(t)}.$$

We estimate the parameter q with a grid from 0 to 1.

The other model we estimate is *modified fictitious play with experimentation* (MFP). Let $o^0_{ij}(t)$ 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 $q_{ij}(t)$,

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

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

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

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, Mailath and Rob (1993) and Young (1993).

As can be easily seen from table 7, the benchmark model performs better (has a lower QDM) than either of the learning models that we estimate. Note, however that the benchmark has more parameters (one per individual) than the others, and that the learning by reinforcement model has only slightly higher QDM than the benchmark (in fact, in the first session the QDM of reinforcement is lower than for the benchmark). In contrast, the modified fictitious play has a quite higher QDM than even the reinforcement model, which has less parameters, for all sessions. This results reproduce quite closely the ones of Chen and Tang (1996) 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.

7 Conclusions

In this paper we have used the risk dominance criterion to select among multiple equilibria in a vertical differentiation model with two firms. In the theoretical analysis, it turns out that risk dominance selects the equilibrium where the high quality is produced by the firm which has some advantage.

To test for the plausibility of this equilibrium selection, we gathered experimental evidence on a coordination game that reproduces the features of the economic model we analyze. 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 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, Mailath and Rob (1993) or Young (1993). In the game we have studied the two approaches coincide but this need not be the case in general. If the two theories 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.

References

BINMORE, K. (1992). *Fun and Games.* D.C. Heath, Lexington MA.

BINMORE, K., A. SHAKED, AND J. SUTTON (1989), “An outside option experiment, ” *Quarterly Journal of Economics*, 104:753–770.

BÖRGERS, T. AND R. SARIN (1993). “Learning throught reinforcement and replicator dynamics, ” Discussion Paper 93–13, UCL.

BREZIS, E., P.R. KRUGMAN, AND D. TSIDDON (1993), “Leapfrogging in international competition: A theory of cycles in national technological leadership, ” *American Economic Review*, 83:1211–1219.

BROSETA, B. (1993). “Strategic uncertainty and learning in coordination games, ” WP 93–34, UC San Diego.

BUDD, C., C. HARRIS, AND J. VICKERS (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. AND M. MOTTA (1996). “Country asymmetries, endogenous product choice and the speed of trade liberalization, ” WP 1326, CEPR, London.

CARLSSON, H. AND E. VAN DAMME (1993), “Global games and equilibrium selection, ” *Econometrica*, 61:989–1018.

CHEN, Y. AND F.F. TANG (1996). “Learning and incentive compatible mechanisms for public goods provision: An experimental study, ” Discussion Paper A-529, Universität Bonn. SFB 303.

COOPER, R., D.V. DEJONG, R. FORSYTHE, AND T.W. ROSS (1989), “Communication in the battle of sexes games: Some experimental results, ” *RAND Journal of Economics*, 20:568–57.

——— (1993), “Forward induction in the battle of sexes game, ” *American Economic Review*, 83:1303–1316.

CRAWFORD, V. (1995), “Adaptive dymaics in coordination games, ” *Econometrica*, 63:103–143.

DASGUPTA, P. AND J. STIGLITZ (1988), “Learning–by–doing, market structure and industrial and trade policy, ” *Oxford Economic Papers*, 40:246–268.

DAVIS, D. AND C.A. HOLT (1993). *Experimental Economics*. Princeton University Press, Princeton NJ.

FARRELL, J. (1987), “Cheap talk, coordination and entry, ” *Rand Journal of Economics*, 18(34–39).

FARRELL, J. AND C. SHAPIRO (1988), “Dynamic competition with switching costs, ” *RAND Journal of Economics*, 19:123–137.

FLAM, H. AND E. HELPMAN (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.

FRIEDMAN, D. AND S. SUNDER (1994). *Experimental Methods: A Primer for Economists*. Princeton University Press, Princeton NJ.

FUDENBERG, D. AND C. HARRIS (1992), “Evolutionary dynamics in games with aggregate shocks, ” *Journal of Economic Theory*, 57:420–442.

GIBBONS, R. (1992). *Game Theory for Applied Economists*. Princeton University Press, Princeton NJ.

GLÜTZ, W., R. SCHMITTBERGER, AND B. SCHWARZE (1994), “An experimental analysis of ultimatum bargaining, ” *Journal of Economic Behaviour and Organization*, 3:367–388.

GROSSMAN, G. AND E. HELPMAN (1991), “Quality ladders and product cycles, ” *Quarterly Journal of Economics*, 106:557–586.

GUYER, M. AND A. RAPOPORT (1972), “2×2 games played once, ” *Journal of Conflict Resolution*, 62:409–31.

HARSANYI, J. (1995), “A new thoery of equilibrium selection for games with complete information, ” *Games and Economic Behavior*, 8:91–122.

HARSANYI, J. AND R. SELTEN (1988). *A General Theory of Equilibrium Selection in Games*. The MIT Press, Cambridge, MA.

HOLT, C. (1995). “Industrial organization: A survey of laboratory research, ” in Kagel, J. and A.E. Roth, editors, *Handbook of Experimental Economics*. Cambridge University Press, New York, NY.

KAGEL, J. AND A.E. ROTH (1995). *Handbook of Experimental Economics*. Cambridge University Press, New York, NY.

KANDORI, M., G.J. MAILATH, AND R. ROB (1993), “Learning, mutation and long–run equilibirum in games, ” *Econometrica*, 61:29–56.

MOTTA, M. (1993), “Endogenous quality choice: Price vs. quatity competition, ” *Journal of Industrial Economics*, 41:113–131.

MOTTA, M., J.F. THISSE, AND A. CABRALES (1997). “On the persistence of leadership and leapfrogging in international trade, ” Discussion Paper 1106, CEPR. Forthcoming in the *International Economic Review*.

MUSSA, M. AND S. ROSEN (1978), “Monopoly and product quality, ” *Journal of Economic Theory*, 18:301–317.

NELSON, R. AND S.G. WINTER (1985). *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.

PORTER, M. (1990). *The Competitive Advantage of Nations*. McMilan, London.

REINGANUM, J. (1985), “Innovation and industry evolution, ” *Quarterly Journal of Economics*, 100:81–99.

ROTH, A. AND I. EREV (1995), “Learning in extensive games: Experimental data and simple dynamic models in the intermediate term, ” *Games and Economic Behavior*, 8:164–212.

ROTH, A. AND M.W.K. MALOUF (1979), “Game-theoretic models and the role of bargaining, ” *Psychological Review*, 86:574–594.

SACCO, P. (1996), “Coordinating coordination failures: Does risk–dominance help?, ” *Economic Notes*, 25:369–392.

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., A. SADRIEH, AND K. ABBINK (1995). “Money does not induce risk neutral behavior, but binary lotteries do even worse, ” Discussion Paper B-343, Universität Bonn, SFB 303.

SHAKED, A. AND J. SUTTON (1982), “Relaxing price competition through product differentiation, ” *Review of Economic Studies*, 49:3–13.

SUTTON, J. (1991). *Sunk Costs and Market Structure*. The MIT Press, Cambridge MA.

TANG, F. (1996). “Anticipatory learning in two–person games: An experimental study; Part II: Learning, ” Discussion Paper B–362, Universität Bonn. SFB 303.

VAN HUYCK, J., R. BATTALIO, AND R. BEIL (1990), “Tacit coordination games, strategic uncertainty and coordination failure, ” *American Economic Review*, 80:2234–249.

YOUNG, P. (1993), “The evolution of conventions, ” *Econometrica*, 61:57–84.

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[13]. 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 hands 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[14]

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

[13] Spanish currency

[14] This paragraph obviously did not appear in the sessions where we did not correct for risk aversion

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 you 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 you couple, as indicated in the following example:

Table 5:

		Choice of B	
		0	1
Choice	0	292, 164	612, 362
of *A*	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 you 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 point 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 5, 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 then 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 6, 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.

Table 1:

Matrix 1:

		Choice of *B* 0	1
Choice	0	29, 29	594, 365
of *A*	1	365, 594	338, 338

Matrix 2:

		Choice of *B* 0	1
Choice	0	123, 78	600, 364
of *A*	1	365, 567	348, 341

Matrix 3:

		Choice of *B* 0	1
Choice	0	193, 114	604, 363
of *A*	1	365, 545	356, 343

Matrix 4:

		Choice of *B* 0	1
Choice	0	248, 142	608, 362
of *A*	1	364, 526	362, 345

Matrix 5:

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

Matrix 6:

		Choice of *B* 0	1
Choice	0	327, 181	615, 361
of *A*	1	364, 496	371, 347

Table 2:

$k = 1.1$

t	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
$u_A u_B$	25.28 4.49	25.28 4.49	25.28 4.49	25.28 4.49	25.28 4.49	25.28 4.49	25.28 4.49	25.28 4.49	25.28 4.49	25.28 4.49	25.28 4.49
			4.70 23.08	4.70 23.08	4.70 23.08	4.70 23.08	4.70 23.08	4.70 23.08	4.70 23.08	4.70 23.08	4.70 23.08

$k = 1.2$

t	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
$u_A u_B$	25.25 4.20	25.25 4.20	25.25 4.20	25.25 4.20	25.25 4.20	25.25 4.20	25.25 4.20	25.25 4.20	25.25 4.20	25.25 4.20	25.25 4.20
				4.59 20.65							
					4.59 21.03	4.59 21.03					
							4.59 21.20	4.59 21.20	4.59 21.20	4.59 21.20	4.59 21.20

$k = 1.3$

t	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
$u_A u_B$	25.22 3.94	25.22 3.94	25.22 3.94	25.22 3.94	25.22 3.94	25.22 3.94	25.22 3.94	25.22 3.94	25.22 3.94	25.22 3.94	25.22 3.94
						4.48 17.41					
							4.48 18.19				
								4.48 18.67			
									4.48 19.14		
										4.48 19.45	
											4.48 19.61

$k = 1.4$

t	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
$u_A u_B$	25.19 3.71	25.19 3.71	25.19 3.71	25.19 3.71	25.19 3.71	25.19 3.71	25.19 3.71	25.19 3.71	25.19 3.71	25.19 3.71	25.19 3.71
								4.37 16.79			
									4.37 17.37		
										4.37 17.81	
											4.37 18.25

$k = 1.5$

t	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
$u_A u_B$	25.17 3.50	25.17 3.50	25.17 3.50	25.17 3.50	25.17 3.50	25.17 3.50	25.17 3.50	25.17 3.50	25.17 3.50	25.17 3.50	25.17 3.50
									4.48 15.98		
										4.27 16.52	
											4.27 17.07

Table 3:

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

Table 4:

Matrix	% of (0,1)	% of (1,0)	% of (0,0)	% of (1,1)
All	41.6	17.5	27.7	13.2
1	15.5	41.0	16.8	26.7
2	26.2	28.8	31.9	13.1
3	40.0	11.1	32.5	16.4
4	45.9	11.2	31.7	11.2
5	43.7	16.3	26.7	13.3
6	70.8	2.2	21.2	5.8

Table 5: 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	P-value
Trial	24	0.17887
1-5	60	0.00029
6-10	60	0.00001
11-15	60	0.00037
16-20	60	0.00183
21-25	58	0.00301
26-30	29	0.00000

Table 6: Test for a linear trend increasing with the degree of asymmetry

	Mean Difference							
Periods	Matrix 1	Matrix 2	Matrix 3	Matrix 4	Matrix 5	Matrix 6	Raw P-Value	Adjusted P-Value
Trial	0.50	-1.00	2.25	0.00	0.67	3.00	0.1075	0.3000
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

Table 7:

Session Number	Matrix	QDM(B)	QDM(LR)	$q(LR)$	QDM(MFP)	$q(MFP)$	$\delta(MFP)$
1	5	10.3	9.96	0.7	11.12	0.55	0.34
2	2	5.7	6.3	0.5	8.42	1	0.22
3	6	6.5	7.52	1	9.69	1	0.26
4	3	6.8	7.76	0.7	11.30	0.75	0.44
5	1	9.5	10.55	0.95	12.26	0	0.44
6	4	8.9	9.17	0.85	11.70	0.90	0.38
7	5	8.1	8.28	0.75	11.96	0.80	0.28
8	2	10.0	10.68	0.8	14.23	0.65	0.44
9	6	4.7	5.62	0.9	9.86	0.90	0.20
10	5	9.4	10.55	0.95	14.54	0.90	0.42
11	2	12.2	13.47	1	14.89	0.70	0.46
12	3	11.4	12.76	0.95	14.94	0.90	0.46

Figure 1: Qualities and profits at the candidate equilibria as a function of the cost efficiency gap, k. Left hand side: qualities (top) and profits (bottom) when firm A is the high quality firm. Right hand side: qualities (top) and profits (bottom) when firm B is the high quality firm.

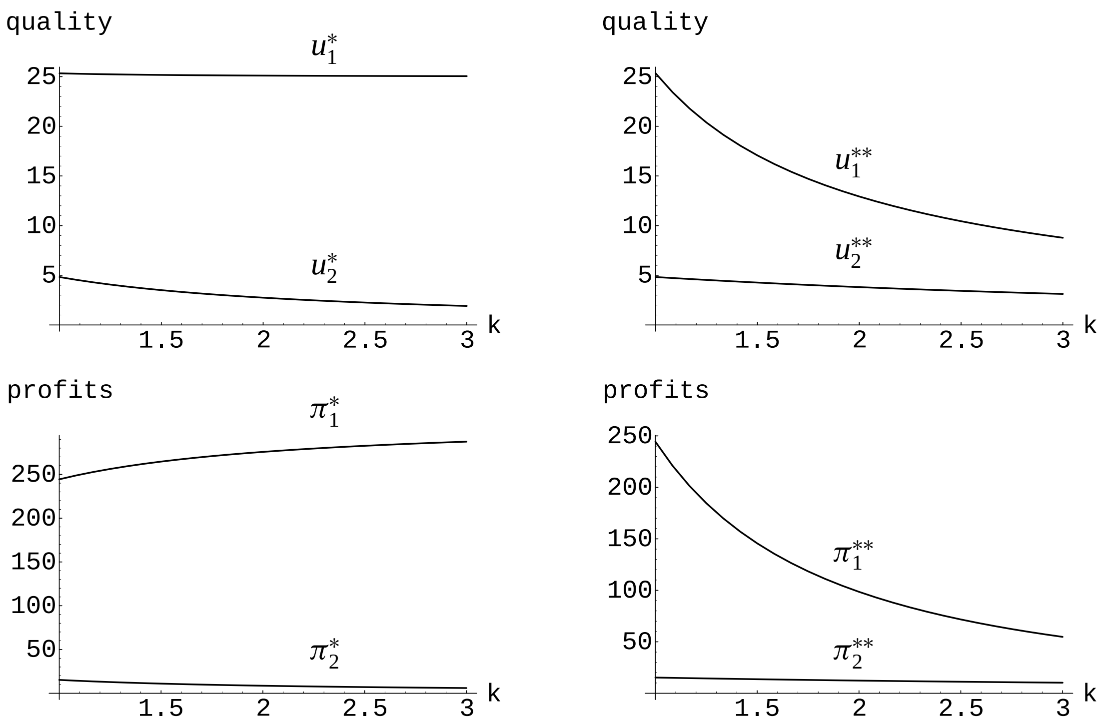

Figure 2: Profits of firm A, at the candidate equilibrium where it produces the low quality, and when it optimally deviates from it to produce a higher quality

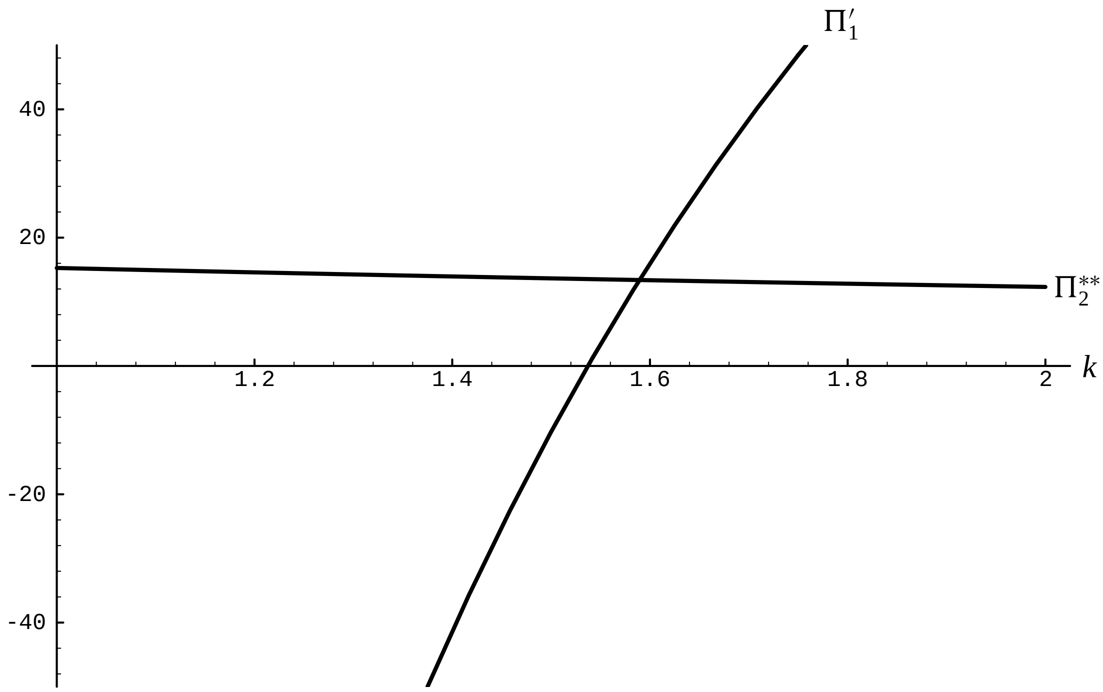

Figure 3: Gain functions for equilibrium predictions

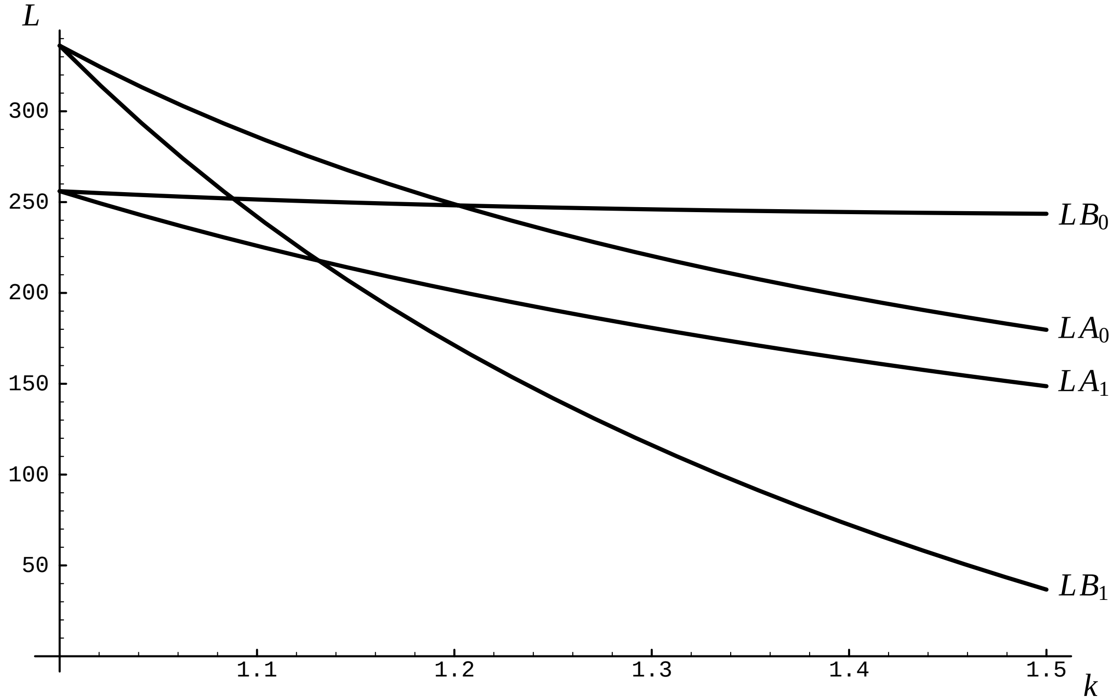

Figure 4:

Difference between equilibria: all experiments

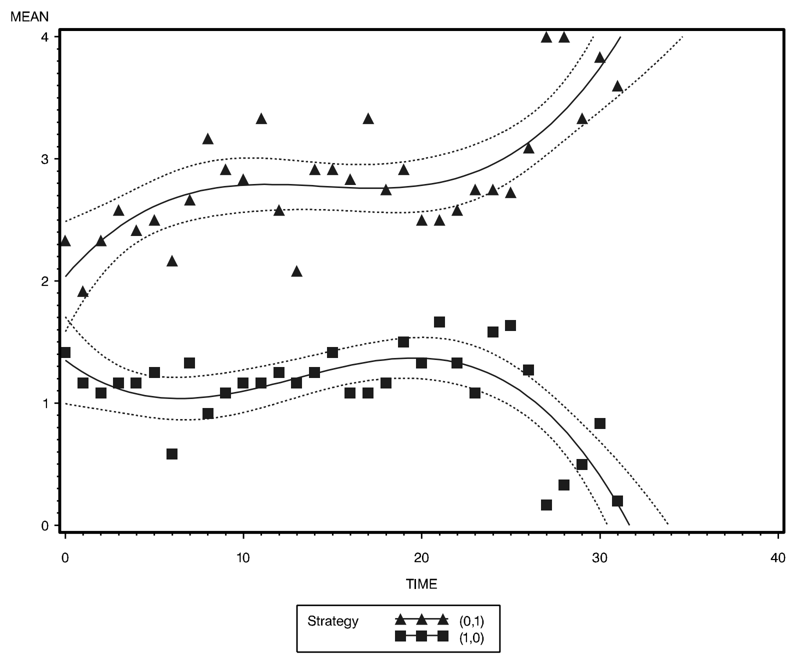

Note: 95 % confidence intervals around the mean (– – –)

Figure 5: Difference in Equilibria: Original Experiments

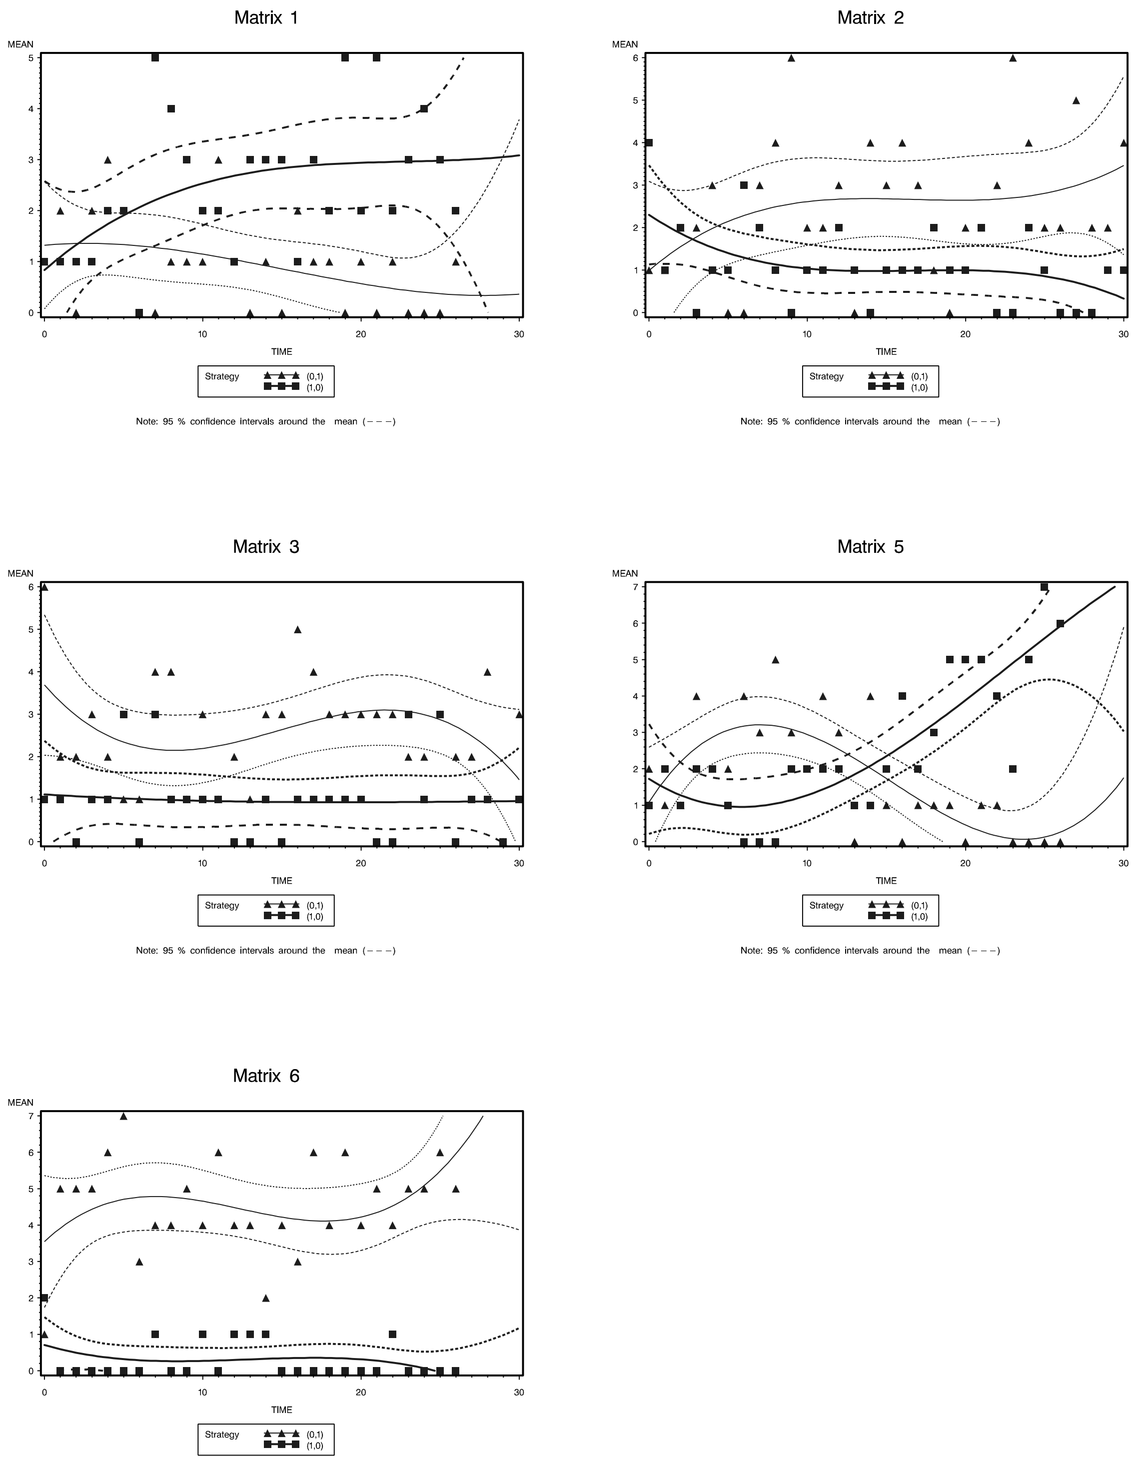

Figure 6: Difference in Equilibria: Replications

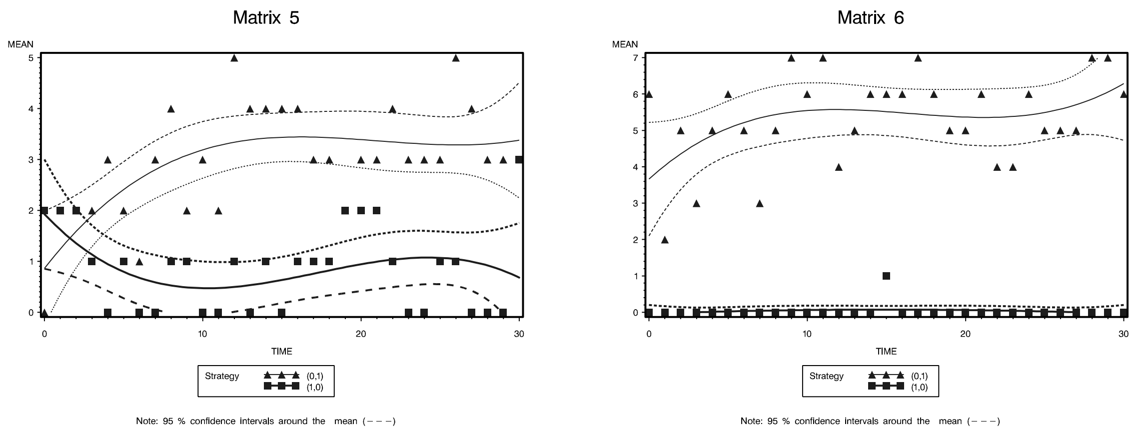

Figure 7: Difference in Equilibria: Experience Effects

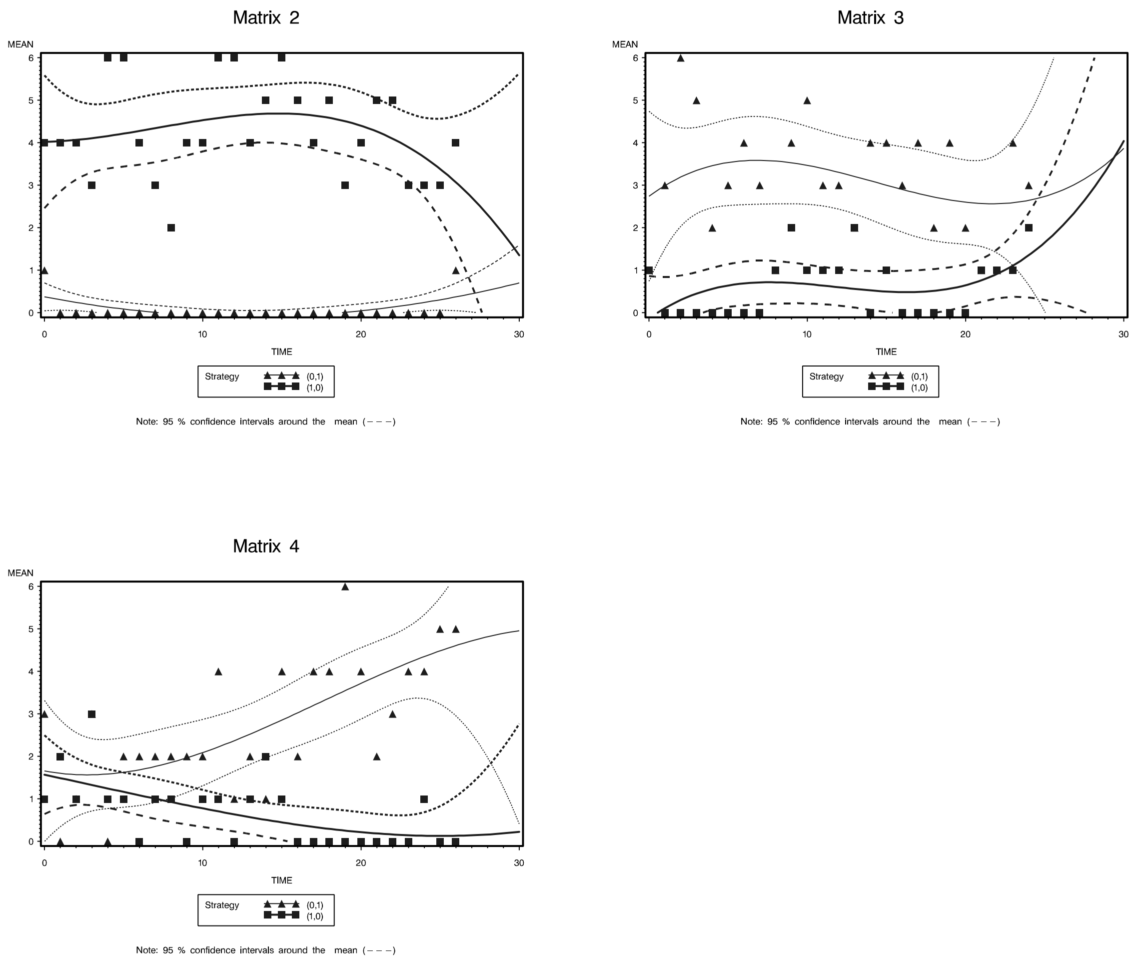

Figure 8: Difference in Equilibria: Lottery Effects