Università degli Studi di Bari
Dipartimento di Scienze Economiche
Quaderni del Dipartimento

Southern Europe Research in Economic Studies – SERIES

SERIES N. 8

A Dynamic Game of Technology Diffusion under an Emission Trading Regulation: A Pilot Experiment

Ivana capozza

A Dynamic Game of Technology Diffusion under an Emission Trading Regulation: A Pilot Experiment.

Ivana Capozza*

April, 2003

Abstract

In this paper we investigate how the interaction between the product and the emission permit markets may affect firms' propensity to adopt cleaner technologies. The adoption of a cleaner technology has the direct effect of reducing the compliance cost of the firm, but it also involves a strategic decision, if the industry is not perfectly competitive. We look at this problem from both a theoretical and an experimental point of view. We develop a model of duopoly, in which two firms engage in quantity competition in the output market and behave as price takers in the permit market. Firms have the possibility of investing in a cleaner production technology, which is available on the market at some cost. We set up a dynamic game over an infinite horizon in order to investigate firms' investment decisions: in each period, each firm decides whether to invest in the new technology or not. The stationary equilibria to this game crucially depend on both the cost of switching to the cleanest technology and the emission cap. Technology diffusion is one of the possible equilibria of the game. In order to test the predictions of the theory, we design and implement an "innovation experiment" that replicates the "innovation game". The results of our pilot experiment suggest that firms' behaviour will eventually lead to innovation diffusion.

Key Words: tradable permits, technology adoption, oligopoly, laboratory experiments.
JEL Categories: C91, L13, O30, Q28

1 Introduction

One central concern of environmental policy is how it can stimulate innovation and diffusion of cleaner technologies. Alternative environmental policy instruments can have

* Dipartimento di Scienze Economiche, Università degli Studi di Bari, Via Camillo Rosalba 53, 70123 Bari (Italy). Email: icapozza@yahoo.com

I would like to thank John Hey for his constant and precious guide and Andrea Morone for his helpful suggestions and for the stimulating discussions I had with him. All remaining errors are solely my own responsibility.

significantly different effects on the rate and direction of technological change. Most of the literature concerning the effect of environmental policy on cleaner technology innovation and diffusion has focused on comparing different policy approaches. There is an almost general agreement that market-based instruments are superior to command-and-control regulation in this respect (Denicolò, 1999; Downing and White, 1986; Fisher et al. 1998, Jung et al., 1996; Milliman and Prince, 1989), though these studies provide different conclusions about which of these market-based policies is the most effective in inducing an "environmental technological change". Nonetheless, emission trading is claimed to be one of the most effective instruments for pollution control, both in a static context and in the long run. The theoretical arguments have led to a growing interest in the possible applications of this policy, both at national and international levels; this interest is also due to some successful applications, mainly in the US.

Nonetheless, only a few studies have dealt with how an emission trading scheme has to be designed in order to induce innovation[1] (Laffont and Tirole, 1996). *What are the variables that affect firms' technology choice under an emission trading scheme?* In this work, we would try to answer this question, focusing on the variables that the environmental regulator may adjust in order to incentive polluters to shift to a cleaner technology.

Moreover, there is little rigorous evidence concerning the ability of tradable permits to incentive innovation and adoption of cleaner technologies, mostly because of the scarcity of available data (Jaffe et al., 2002). This observation has led us to adopt the experimental approach to analyse firms' technology choice. *Do laboratory subjects decide to innovate as theory predicts?*

As it has been pointed out in some studies (Fershtman and deZeew, 1995; Montero, 2002; Requate, 1998), considering only the allowance market may lead to biased conclusions, since the adoption of a cleaner technology has a direct cost-reducing effect internal to the firm, but also implies strategic effects arising from the interaction between the output and the permit markets. For this reason, we depart from the competitive assumption, explicitly considering the production decision in a Cournot duopoly, where the technology is modelled in terms of emissions per unit of output, rather than in terms of abatement cost. We focus on the effect of the interaction between the output and the permit markets on the diffusion of the new technology, i.e. we look at the conditions

[1] For example, Laffont and Tirole (1996a and 1996b) claim that future markets of permits provide higher incentives to innovate than do spot markets.

under which one or both firms are induced to invest in a new technology and thus on what it is possible to do to make diffusion more likely.

In order to analyse the incentives to innovate in the long run, we set up a dynamic game: in each period, each duopolist has to decide whether to invest in the new production process or not, on the basis of the profits they can earn during their whole life, which we assume to be infinite. We solve the game looking for a stationary symmetric equilibrium, which may involve mixed strategies. We find that this game leads to multiple stationary equilibria, which crucially depend on the cost of switching to the cleanest technology and on the emission cap. In particular, one of these possible outcomes consists in both firms adopting the most efficient technology (*diffusion outcome*). Nonetheless, it is not possible to predict which outcome will actually occur. An experimental investigation can help us in seeing this. We design and implement an "innovation experiment" that replicates the "innovation game". Our aim is to see whether subjects tend to converge to a specific equilibrium among those that have been identified in the theoretical analysis.

The rest of this paper is organised as follows. In the next section we develop a model of duopoly, and we determine how the permit price and the per-period profits change when firms move to a cleaner technology. In section 3, we model firms' strategic decision of adopting a cleaner production technology as a dynamic "innovation" game. Section 4 describes the experiment that was implemented to test the predictions of the theoretical model. The final section concludes, drawing some lessons for environmental policy design.

2 The Model

We consider two profit-maximising firms producing a homogeneous good and competing à la Cournot in the output market. Let y_i be firm i's output level, for $i = 1,2$. The inverse demand for output is linear $P(y_1 + y_2) = a - b(y_1 + y_2)$ and firms have identical constant marginal production cost c, with $c < a$.

The production process generates pollution according to a constant emission-output ratio k_j, for $j = 1,2$, such that firm i's emission level is $e_i = k_i y_i$. There are only two possible technologies in this setting, k_1 and k_2, with $k_j \geq 0$ for $j = 1,2$[2], and $k_1 \leq k_2 \leq 2k_1$[3]; hence, k_1 is the cleanest technology[4]. The technology is the only thing for which firms may differ one another: starting from an identical technology k_2, each firm may decide to switch to the other and more efficient technology, bearing a sunk cost C.

The government implements a market of emission permits in order to control pollution by firms. Because of this regulation, firm i must hold one permit to discharge one unit of emissions. There is no possibility of abatement in this model: given the adopted technology k_j, the only way to reduce pollution is to cut production.

The emission cap is fixed at E and firms do not have any initial endowments of licences: both of them are buyers on the allowance market.

We assume discrete time and an infinite horizon: firms operate in both markets for an infinite number of periods. In each period, firms face the following decision problem: given the permit price[5], firms engage in Cournot competition. The quantity choice implies a predetermined permit requirement, which depends upon the adopted technology. Hence, in each period and taking as given the permit price, each firm determines the demand for permits needed to produce its profit-maximising quantity of the good. The equilibrium allowance price, q, is derived endogenously by imposing a market clearing condition[6], given the fixed permit supply and the factor demands for permits implied by the state of technology. Finally, output and profit levels are resolved. Aggregate emissions are fixed by definition at the level allowed by the total supply of licences.

In the initial period, firms use the same production technology and an emission reducing technology becomes available for purchase. In each period, firm i may decide to pay a sunk cost C, investing in this new technology. Once a firm has adopted the cleanest available technology, it cannot return to the previous one, neither it can innovate further:

[2] We assume that the constant marginal pollution is strictly positive, i.e. it is not possible to produce any unit of the commodity without polluting.

[3] This condition is imposed in order to assure existence and uniqueness of equilibrium. If this condition is not met, the demand for permit of the most efficient firm may be upward sloping.

[4] Note that the subscript j does not necessarily mean that firm i adopts technology $j = i$.

[5] Firms may have expectations upon permit price.

[6] We can think of an auctioneer who sells the permit at the market clearing price: given a price (previous period price or a price announced by the auctioneer), firms believe they can purchase any amount of permits at that price and announce the quantities they will put on the output market; the auctioneer anticipates the factor demands for permits and announces the market clearing price (Requate, 1998).

hence, investment is irreversible and firms can innovate[7] only once in their life. Therefore, in each period firms can face one of the following situations:

- State 0: they both use the least efficient technology k_2, i.e. none of them has innovated;
- State 1: only one of them has switched to the most efficient technology k_1; we assume that the innovating firm is firm 1, whereas firm 2 is the follower;
- State 2: both firms have adopted the best emission-output ratio k_1.

We first analyse the firms' profit maximising decisions in each period and for each possible combination of production technologies, disregarding the technology choice. Then, we analyse firms' investment decisions. In what follows the superscript $s = 0,1,2$ refers to the state, the subscript $i = 1,2$ refers to the firm and the subscript $j = 1,2$ refers to the emission-output ratio.

When both firms use the same technology, they are identical in all respects and the solution of the model is analogous to the standard Cournot duopoly case. Let us consider state 0 (the analysis of state 2 is analogous). Each firm maximises its profits w.r.t. output y_i for $i = 1,2$:

$$\max_{y_i}[a - b(y_i + y_h)]y_i - cy_i - qk_2 y_i$$

Solving both firms' maximization problem yields the equilibrium output levels[8], the demand for permits (multiplying the optimal output quantities by firms' technological parameter), and the equilibrium permit price. In state 0, the equilibrium permit price[9] will be positive only if $0 < E < \alpha^0$, with $\alpha^0 = [2k_2(a-c)]/3b$; in state 2, the equilibrium permit price will be positive only if $0 < E < \alpha^2$, with $\alpha^2 = [2k_1(a-c)]/3b$ and $\alpha^2 < \alpha^0$: there are emission caps for which the adoption of a cleaner technology by both firms

[7] We use the terms "innovation", "adoption" and "investment" as equivalent, meaning that a firm change its technology paying a given cost. More rigorously, the term "innovation" is used to indicate the result of a research and development process, whereas "adoption" is used to denote the switch to a new technology which is already available on the market and can be used at some cost. (Tirole, 1998)

[8] The equilibrium output levels are identical, and equal to $y_1^0(q) = y_2^0(q) = (a - c - qk_2)/3b$, where the term qk_2 is an extra bit of the marginal production cost. If everything stays equal, reducing the emission-output ratio from k_2 to k_1 makes both firms' output levels increase.

[9] The equilibrium permit price in state 0 is $q^0 = [2k_2(a-c) - 3Eb]/2k_2^2$. The equilibrium permit price in state 2 is $q^2 = [2k_1(a-c) - 3Eb]/2k_1^2$.

drives the initially positive price to zero[10]. For the moment we make the following assumption:

Assumption 1. The emission cap is such that $0 < E < \alpha^2$.

Assumption 1 implies that in case of both firms adopting k_1 the permit price is strictly positive.

Proposition 1. Under assumption 1, when both firms adopt the cleanest emission-output ratio k_1, we have:

- if $\beta^2 < E < \alpha^2$, the equilibrium allowance price is lower than in state 0;
- if $0 < E < \beta^2$, the equilibrium allowance price is higher than in state 0;

where $\beta^2 = 2k_1k_2(a-c)/3b(k_1+k_2) > 0$.

Proposition 2. Under assumption 1, when both firms adopt the cleanest emission-output ratio k_1, it is $\pi^2 > \pi^0$: firms' (output and) profits are higher than when both firms adopt the least efficient technology[11].

Let us now consider the case in which only one firm has innovated (state 1). We assume that firm 1 has adopted the lowest emission-output ratio, k_1, whereas the other firm produces discharging $k_2 > k_1$ units of emission per unit of output. Solving firms' maximisation problem gives their output functions[12]. From the optimal output functions we can derive firm 1's and 2's demands for permits.

Proposition 3. Let us denote $\theta = [k_1(k_2-k_1)(a-c)]/(2k_2-k_1)b > 0$. If $E \leq \theta$, the innovating firm becomes a monopolist on the output market, the other firm makes zero profit and is forced to exit (*drastic* innovation), otherwise both firms stay on the market (*nondrastic* innovation) and make strictly positive profits.

---

[10] Notice that if the permit price is null, both firms produce an output quantity equal to $(a-c)/3b$. This is the standard result of a Cournot game with identical firms and linear output demand. It is the levels of output each firm would produce if there were no environmental regulation (Business As Usual, BAU). α^0 and α^2 are the level of emissions associated to the BAU level of output, when the emission-output ratio is k_2 and k_1, respectively.

[11] Firms' identical equilibrium profits are $\pi^0 = b(E/2k_2)^2$ and $\pi^2 = b(E/2k_1)^2$, when both firms use the old and the new technology, respectively. Profits are constrained by the environmental regulation.

[12] Firm 1's and firm'2 output functions (optimal output levels as function of the permit price q) in state 1 are $y_1^1(q) = [a-c-q(2k_1-k_2)]/3b$ and $y_2^1(q) = [a-c-q(2k_2-k_1)]/3b$, respectively.

We focus on the case of *nondrastic* innovation. The equilibrium permit price[13] will be positive for $E < \alpha^1$, with $\alpha^1 = [(k_1 + k_2)(a - c)]/3b > 0$ and $\alpha^2 < \alpha^1 < \alpha^0$. Let us make the following assumption:

Assumption 2. The emission cap is such that $\theta < E < \alpha^2$.

Assumption 2 will be maintained throughout the analysis. Note that assumption 2 is stronger than assumption 1.

Proposition 4. Under assumption 2, when only one firm adopts the cleanest emission-output ratio k_1, we have

- if $\frac{1}{2}k_2 < k_1 \le \frac{2}{3}k_2$, it is $q^1 < q^0$: the equilibrium allowance price is lower than in state 0 whichever the emission cap is;
- if $k_1 > \frac{2}{3}k_2$ the equilibrium allowance price can be lower or higher than in state 0:
    - it is $q^1 < q^0$ for $\beta^1 < E < \alpha^2$;
    - it is $q^1 > q^0$ for $\theta < E < \beta^1$;

with $\beta^1 = k_2(2k_1 - k_2)(a - c)/3bk_1 > 0$ and $\beta^1 < \beta^2$.

Proposition 5. Under assumption 2,

- if $\theta < E < \varepsilon$, it is $\pi_2^1 < \pi^0$ and $\pi^2 < \pi_1^1$: innovation by one firm hurts the other firm, whether the firm that adopts k_1 is the first or the second one to do so;
- if $\varepsilon < E < \alpha^2$, it is $\pi_2^1 > \pi^0$ and $\pi^2 > \pi_1^1$: innovation by one firm benefits the other firm whether the firm that adopts k_1 is the first or the second one to do so;

where $\varepsilon = k_1 k_2 (a - c)/(k_2 + k_1)b > 0$ and $\beta^1 < \varepsilon < \alpha^2$.

Proposition 6. From assumption 2 and propositions 2 and 5, there are two possible ranking of per-period profits:

- if $E < \varepsilon$ then $\pi_1^1 > \pi^2 > \pi^0 > \pi_2^1$;
- if $E > \varepsilon$ then $\pi^2 > \pi_1^1 > \pi_2^1 > \pi^0$.

[13] The equilibrium permit price in state 1 is $q^1 = [(a - c)(k_1 - k_2) - 3Eb]/2(k_1^2 - k_1 k_2 - k_2^2)$.

Proposition 7. Under assumption 2,

- if $E < \varepsilon$, it is $(\pi_1^1 - \pi^0) > (\pi^2 - \pi_2^1)$;
- if $E > \varepsilon$, it is $(\pi_1^1 - \pi^0) < (\pi^2 - \pi_2^1)$.

In general, taking the permit price as fixed, the adoption of a cleaner technology makes it more convenient for a firm to increase its output level. Whether the firm is the leader or the follower or both firms innovate simultaneously, innovation cuts overall production costs by reducing permit requirement and expenditure on licence purchase per unit of output: given the price, a reduction in the emission-output ratio is equivalent to a decrease in the marginal cost. Hence, when only one firms adopts the new technology, innovation enables it to increase its market share and profits at the expense of its rival. Despite the higher average abatement by the innovating firm, the allowance price may or may not decrease, depending on the emission cap and on the technology combination (k_1, k_2): the smaller the permit supply is, the bigger the innovation effort must be in order to determine a sufficient decrease in permit requirement and thus a fall in the permit price. Consequently, the non-innovating firm may be advantaged by the other firm's investment: the lower licence price may enable it to increase its production and profits, without investing in a cleaner technology. However, firm 2 can exploit the price drop only if this drop is sufficiently large. This happens for $E > \varepsilon$. On the other hand, if firm 2 follows and adopts the new technology (state 2), its output level is always higher than in state 1, whereas firm 1's production may go up or down. As for the non-innovating firm in state 1, when imitation[14] occurs the firm that has innovated first may benefit from imitation and may increase its output (and profits) in state 2 if the permit price decreases sufficiently. The condition for this to occur is again $E > \varepsilon$. Finally, if $E < \varepsilon$, the increment in per-period profits a firm can get by adopting first is higher than the increment in profits it can get by being second. The opposite occurs if the permit supply is sufficiently large.

Till now we have considered the output decisions of firms given their technology. In the following section we consider firms' decisions of investing in a cleaner technology. This strategic decision depends upon the ranking between per-period profits.

[14] We will say that a firm imitates when it is the second to adopt the new technology, i.e. when it is the follower.

3 The game

In analysing the strategic investment decision, we assume discrete time and an infinite horizon. Let us consider the following game, with two players, firm 1 and firm 2, choosing between two possible actions in each period, "*innovate*" and "*not innovate*". In every period t, each firm has the possibility of adopting the best available technology k_1 if it has not already done so, incurring in a once-for-all cost C. If neither firm has innovated at time t, each one can still decide to do so in the next period, facing the same decision problem at time $t+1$.

Hence the game ends when both firms have changed their technology, either because they decided to do so at the same time or because one firm has innovated at a given time and the other has chosen to follow[15]. Since each decision implies a flow of profits, in making their choice in each period, firms must consider this profit flow over their infinite life, discounted with a discount factor ρ, $0<\rho<1$. Firm i's payoff associated to a given combination of firms' actions is its respective lifetime profit as viewed from the beginning of the period[16].

If both firms decide to adopt technology k_1, they get a profit flow of π^2 for all the remaining periods, net of the investment cost:

$$\Pi^2 = \frac{\pi^2}{1-\rho} - C \tag{1}$$

If one firm has innovated (firm 1), the other one (firm 2) must decide whether to follow or not in the next period. Firm 2's action is chosen solving a single firm optimisation problem. Considering two arbitrary periods, t and $t+1$, firm 2 will adopt the best technology in period t rather than in the next one if it gets higher lifetime profits by doing so, i.e. if

$$\pi_2^1 \sum_{z=0}^{t-1} \rho^z + \rho^t \left(\frac{\pi^2}{1-\rho} - C \right) > \pi_2^1 \sum_{z=0}^{t} \rho^z + \rho^{t+1} \left(\frac{\pi^2}{1-\rho} - C \right)$$

---

[15] The game ends since in this case players have no strategic decision to take, but they continue to get their per-period profits for the rest of their life.

[16] Both firms know that if they will be playing the game at an arbitrary time t, it would be because they will not have innovated before time t. Hence they both will have accumulated $\pi^0 \sum_{z=0}^{t-2} \rho^z$ and they will keep this profit for the periods thereafter, whichever decision they will take at time t. This part of firms' payoff will not affect their choice between innovating or not.

which is true for

$$(1-\rho)C < \pi^2 - \pi_2^1 \qquad (2)$$

Condition (2) means that the extra profit firm 2 would get if it followed in each period t rather than the next one, is higher than the savings in the cost due to one period delay.

Proposition 8. Let us denote $(\pi^2 - \pi_2^1)/(1-\rho)$ as $\overline{C}$. If firm 1 innovates, then the other firm either adopts the new technology immediately next period or never, depending on the investment cost. Two cases are possible:

a) *quick imitation*: if $C < \overline{C}$, firm 2 innovates immediately after the other one has; the lifetime profits of the leader and of the follower are

$$\Pi_1^1 = \pi_1^1 - C + \frac{\rho\pi^2}{1-\rho} \quad (3) \quad \text{and} \quad \Pi_2^1 = \pi_2^1 + \frac{\rho\pi^2}{1-\rho} - \rho C \quad (4)$$

respectively, since state 1 will last only one period and firms pass to state 2;

b) *infinite delay*: if $C > \overline{C}$, firm 2 always delays adoption (never adopts); the lifetime profits of the leader and of the follower are

$$\Pi_1^1 = \frac{\pi_1^1}{1-\rho} - C \quad (5) \quad \text{and} \quad \Pi_2^1 = \frac{\pi_2^1}{1-\rho} \quad (6)$$

respectively, since state 1 will last forever.

If neither firm has innovated at time t, the game is repeated the next period. The number of periods the game will last is potentially infinite: the game stops (but the profit flows continue) if one or both firms innovate, since when only one firm has innovated the other either follows immediately or never. Therefore, in each period the payoff Π^0 that each firm gets if no one innovates, depends on what firms will do in the following period(s).

Since we solved the part of the game involving only one firm's decision, we can represent the "innovation" game as a symmetric game with two identical firms-players and two actions ("*innovate*" and "*not innovate*") in each period, and with payoffs equal to lifetime profits, net of the investment cost. There are two games, one for the "quick imitation" case and one for the "infinite delay" case.

We can solve the game looking at each subgame, i.e. looking at every period in which neither player has already innovated. Our aim is to solve the game for a stationary (i.e. time independent) Nash equilibrium, possibly involving mixed strategies. If such a symmetric stationary equilibrium exists, it is subgame perfect, since all subgames in which both firms can make a choice have the same structure and payoffs.

In order to see whether a stationary equilibrium to this game exists, we start by considering each subgame and analysing the possible relationship between payoffs, implied by the output and permit market equilibria.

In the quick imitation case, it is $\Pi^2 > \Pi_2^1$ by propositions 6 and 8. If it is also $\Pi_1^1 > \Pi^0$, "*not innovate*" is a strictly dominated strategy for both firms, and the unique possible equilibrium is a symmetric Nash equilibrium (*innovate*, *innovate*): both firms innovate straight away in the first period and the game ends. If it is $\Pi_1^1 < \Pi^0$, neither firm has a strictly dominated strategy and there are three equilibria: two symmetric Nash equilibria in pure strategies, in which both firms either innovate straight away or never, and one symmetric equilibrium in mixed strategies.

In the infinite delay case, it is $\Pi^2 < \Pi_2^1$ by propositions 6 and 8. If it is also $\Pi_1^1 < \Pi^0$, "*innovate*" is a strictly dominated strategy for both firms, and the unique possible equilibrium is a symmetric Nash equilibrium (*not innovate*, *not innovate*). If it is $\Pi_1^1 > \Pi^0$, neither firm has a strictly dominated strategy and there are three equilibria: two asymmetric Nash equilibria in pure strategies, in which a firm innovates straight away and the other never does, and one symmetric equilibrium in mixed strategies.

Therefore, a mixed strategy equilibrium is a possible solution to this game in both the quick imitation and the infinite delay cases. Whether or not this equilibrium actually arises depends crucially on the parameters, since they affect both per-period and lifetime profits. We will focus on the effects of the emission cap and of the investment cost on the equilibrium, since these are the parameters that the environmental regulator might adjust in order to speed up the diffusion of cleaner technologies. In order to determine the critical values of these parameters, we will solve the game assuming mixed strategies and then we will see for which values of E and C this assumption is correct.

There is a mixed strategy equilibrium for each subgame if there are a probability p_1 and a probability p_2 such that firm 1's strategy "do not innovate with probability p_1 and innovate with probability $(1-p_1)$ (conditional on not having innovated before)" is the best response to firm 2's strategy "do not innovate with probability p_2 and innovate with probability $(1-p_2)$ (conditional on not having innovated before)", and viceversa. Since firms are identical, it is $p_1 = p_2 = p$.

Let us consider firm i's point of view and assume that the other firm plays in each period the strategy $(p, 1-p)$. If firm i assumes this, it will innovate in a given period t if the expected payoff of doing so is higher than the expected payoff of not innovating, and viceversa. Firm i's expected payoff of innovating and of not innovating are

$$V_I = p\Pi_1^1 + (1-p)\Pi^2 \quad (7) \quad \text{and} \quad V_N = \frac{p\pi^0 + (1-p)\Pi_2^1}{1-\rho p} \quad (8)$$

respectively[17].

Both V_N and V_I are functions of p, and their exact forms depend on whether there is quick imitation or infinite delay. Each firm's optimal mixed strategy is determined solving $V_N = V_I$ for p. Let us denote this solution as p^*. Hence, if a symmetric mixed strategy equilibrium exists, it is such that in any period both firms do not innovate with probability p^* and innovate with probability $(1-p^*)$.

It can be noticed that $V_N = V_I$ yields a quadratic expression in p. We solved this equation for p for both the quick imitation and the infinite delay cases[18]. It can be shown that in both cases, there are values of the investment cost C and of the emission cap E for which a symmetric mixed strategy equilibrium $((p^*, 1-p^*), (p^*, 1-p^*))$ actually exists.

---

[17] V_N is determined considering that:
-if at any time t one firm innovates, which happens with probability $(1-p)$, the other firm gets Π_2^1;
-if one firm does not innovate, which occurs with probability p, the other firm gets π^0 for the current period and the expected value of not innovating for the future periods, discounted with the discount factor ρ.

[18] The solution for the quick imitation case is $p^* = \left[\pi_1^1 - \pi^0 + \pi_2^1 - \pi^2 + \sqrt{(\pi_1^1 - \pi^0 + \pi_2^1 - \pi^2)^2 + 4\rho(\pi^2 - \pi_2^1 - C(1-\rho))(\pi_1^1 - \pi^2)}\right] \big/ 2(\pi_1^1 - \pi^2)\rho$. The second root of the equation $V_N = V_I$ is in the range between 0 and 1 only for $C > \overline{C}$, which implies that the infinite delay lifetime profits should be considered. The solution for the infinite delay case is

$$p^* = \frac{1}{2(\pi_1^1 - \pi^2)\rho}\Big\{(1-\rho)(\rho C - \pi^0) - (1+\rho)\pi^2 + \pi_2^1 + \pi_1^1 + \\ - \sqrt{((1-\rho)(\rho C - \pi^0) - (1+\rho)\pi^2 + \pi_2^1 + \pi_1^1)^2 + 4\rho(\pi_2^1 - \pi^2 + C(1-\rho))(\pi^2 - \pi_1^1)}\Big\}$$

The second root of the equality $V_N = V_I$ is in the range between 0 and 1 only for $C > \widetilde{C}$, which implies that innovation should never occur, as we will see later.

Let us define $\hat{C}$ as the minimum value of the investment cost required to have a mixed strategy equilibrium in the quick imitation case, where $\hat{C} = \left[\pi_1^1 - \pi^0 - \rho\left(\pi_1^1 - \pi^2\right)\right] / (1 - \rho)$.

Proposition 9. In the quick imitation case, the following equilibria arise:

- if $\hat{C} < C < \overline{C}$, there are a symmetric mixed strategy equilibrium $\left(\left(p^*, 1 - p^*\right), \left(p^*, 1 - p^*\right)\right)$ and two symmetric pure strategy Nash equilibria (*not innovate*, *not innovate*) and (*innovate*, *innovate*);
- if $C < \hat{C}$, there is a pure strategy equilibrium (*innovate*, *innovate*) in every period.

Hence, proposition 9 implies that each firm innovates for sure if the benefit of waiting for one period (savings in the investment cost obtained by delaying of one period the adoption of k_1) is lower than the forgone profits of being first (extra-profit of being the first to adopt the new technology, net of the profit decrease caused by imitation). Indeed, for $C < \hat{C}$ the probability is not defined and we can conclude that there is a pure strategy equilibrium (*innovate*, *innovate*)[19]. On the other hand, $\hat{C} < C < \overline{C}$ implies that if the cost savings of waiting is higher than the opportunity cost of being first but lower than the opportunity cost of being second (recall proposition 8), each firm can do better by waiting for the other to invest and investing with one period delay[20]. Therefore, the mixed strategy equilibrium arises when firms are afraid of failing in coordinate themselves on the Pareto dominant outcome (*not innovate*, *not innovate*): the mixed strategy outcome is worse than the pure strategy (*not innovate*, *not innovate*), but it is better than the pure strategy (*innovate*, *innovate*) There is a discontinuity[21] at $C = \hat{C}$, as it is evident from Figure A1.1 (Appendix 1).

[19] For each firm it is $V_N < V_I$ either if the other firm innovates for sure or if it does not. If one firm adopts k_1 with certainty, the other firm's expected payoff of not innovating collapses to Π_2^1, whereas its expected payoff of innovating degenerates to Π^2, with $\Pi^2 > \Pi_2^1$. If one firm does not adopt the new technology for sure, the other firm's expected payoff of not innovating is $\pi^0 / (1 - \rho)$, whereas its expected payoff of innovating degenerates to Π_1^1, with $\Pi_1^1 > \pi^0 / (1 - \rho)$ for $C < \hat{C}$.

[20] Indeed, each firm would be better off if they both continue to use the old technology forever (since it is $\pi^0 / (1 - \rho) > \Pi^2$), but if one adopts k_1 the other one has to follow in order to limit its loss.

[21] For $C = \hat{C}$ the solution to $V_N = V_I$, is $p^* = 1$, implying that for this value of the investment cost there is a pure strategy equilibrium (*not innovate*, *not innovate*) in every period.

Proposition 10. In the quick imitation case, when a mixed strategy equilibrium arises, the probability of not innovating p^* is decreasing in the investment cost C, other things being equal.

Proposition 10 says that it is more likely to adopt the best technology when the investment cost gets higher. This result seems counterintuitive. One possible explanation may stems from the incentive to play a mixed strategy in the quick imitation case: each firm would like to be the only one to adopt the new technology but knows that it will be followed. When the investment cost is higher, both firms have a greater incentive to wait and see whether the other invests: a firm must ("threaten" to) innovate with a higher probability in order to leave its rival indifferent between innovating or not.

Let us now consider how the mixed strategy equilibrium changes with the emission cap E.

Proposition 11. In the quick imitation case, the following equilibria are feasible :

- if $\underline{\varepsilon} < E < \varepsilon$, there are a symmetric mixed strategy equilibrium $\left(\left(p^*, 1-p^*\right), \left(p^*, 1-p^*\right)\right)$ and two pure strategy Nash equilibria (*not innovate*, *not innovate*) and (*innovate*, *innovate*), provided that $\hat{C} < C < \overline{C}$,
- if $E < \underline{\varepsilon}$ or $E > \varepsilon$, there is a pure strategy equilibrium (*innovate*, *innovate*) in every period[22].

First of all, changes in the emission cap affect the mixed strategy equilibrium because they affect the critical cost range. It can be $\hat{C} < C < \overline{C}$ only for $\underline{\varepsilon} < E < \varepsilon$. For E outside of this range, the relative magnitude of per-period profits is such that the net benefit of being first (extra per-period profit, net of the profit decrease caused by imitation) is always greater than the benefit of being second[23].

---

[22] $\underline{\varepsilon} = \frac{k_1 k_2 (k_1 - k_2)(a-c)((\rho - 1)k_2^2 - k_1^2)}{k_2^2((k_2 - k_1)^2 + 2k_1^2)\rho + (k_2 + k_1)(k_1 - k_2)^3}$ and $\underline{\varepsilon}$ is higher or lower than θ depending on the discount factor and on the technology parameters.

[23] Recalling proposition 7:

- if $\underline{\varepsilon} < E < \varepsilon$, the decrease in per-period profit caused by being followed partially compensates the extra-profit of innovating first, so that a firm can get an higher increment of the per-period profit by adopting k_1 with one period delay: there exist investment cost $C < \overline{C}$ such that the cost savings of waiting is higher than the extra-profit of being first but lower than the extra-profit of being second and a mixed strategy equilibrium is feasible;

Proposition 12. In the quick imitation case, when a mixed strategy equilibrium arises, the probability of not innovating p^* is increasing in the emission cap E, other things being equal.

When the permit supply increases, the benefit of adopting the new technology increases slightly more than the expected value of not innovating; hence, the incentive to adopt increases accordingly, even though the best outcome for both firms is still associated to both staying with the old technology. A firm must thus ("promise") not to innovate with a higher probability in order to leave its rival indifferent between innovating or not

We will now consider the infinite delay case. Let us denote $\widetilde{C}$ the maximum value of the investment cost for which a mixed strategy equilibrium arises in the infinite delay case, where $\widetilde{C} = \left(\pi_1^1 - \pi^0\right)/(1-\rho)$.

Proposition 13. In the infinite delay case, the following equilibria may arise:

- if $\overline{C} < C < \widetilde{C}$, there are a symmetric mixed strategy equilibrium $\left(\left(p^*, 1-p^*\right), \left(p^*, 1-p^*\right)\right)$ and two asymmetric pure strategy Nash equilibria (*innovate*, *not innovate*) and (*not innovate*, *innovate*);
- if $C > \widetilde{C}$, there is a pure strategy equilibrium (*not innovate*, *not innovate*)[24].

Hence, proposition 13 implies that each firm does not innovate if the benefit of waiting for one period are higher than the forgone profits of being first[25]. Indeed, for $C > \widetilde{C}$ the probability is not defined and we can conclude that there is a pure strategy equilibrium

---

- if $E < \underline{\varepsilon}$ (the emission cap is sufficiently strict), the extra-profit of innovating first net of the decrease in per-period profit caused by being followed is higher than the extra-profit of being second: for every $C < \overline{C}$ the cost savings of delaying adoption of one period are lower than the extra-profit that can be realised by adopting technology k_1, and firms innovate straight away;
- if $E > \varepsilon$ (the permit supply is sufficiently large), the extra-profit of being second is higher than the extra-profit of being first; however, imitation leads to an increase in the per-period profit of the first adopter, so that being first is still better: as before, for every $C < \overline{C}$ the cost savings of delaying adoption of one period are lower than the extra-profit that can be realised by adopting technology k_1, and firms innovate straight away.

[24] For $C = \overline{C}$, it is $p^* = 0$. However, as for quick imitation, we leave this case undetermined, since $\overline{C}$ is the value of C for which the non-innovating firm should be indifferent between adopting the new technology immediately in the next period or never. For some values of $C > \widetilde{C}$, there are two solutions to the equation $V_N = V_I$ which implies that there is no mixed strategy equilibrium. However, for $C > \widetilde{C}$ it is $V_N(0) > V_I(0)$ and $V_N(1) > V_I(1)$: there is a unique pure strategy equilibrium (*not innovate*, *not innovate*).

[25] In this case, the other firm will not follow, and the innovating firm does not incur in a reduction of its extra-profit in the following period.

(*not innovate*, *not innovate*)[26]. Therefore, the condition $\overline{C} < C < \widetilde{C}$ implies that if the cost savings of waiting is lower than the opportunity cost of being first, provided that the investment cost is sufficiently high to avoid imitation (recall proposition 8), each firm can do better by trying to be the first and only one to adopt the new technology, i.e. by *preempting* its rival. Each firm can find it convenient to play a mixed strategy in order to confuse its rival and to be able to reach the innovator's payoff. The mixed strategy outcome is worse than the pure strategy one (*innovate*, *not innovate*), but it is better than the pure strategy one (*not innovate*, *innovate*). There is a discontinuity at $C = \widetilde{C}$, as it is evident from Figure A1.2.

Proposition 14. In the infinite delay case, when a mixed strategy equilibrium arises, the probability of not innovating p^* is increasing in the investment cost C, other things being equal.

This is a more intuitive result than we have for the quick imitation case. If the investment cost increases, the advantage of being the first to adopt the new technology decreases; hence, both the incentive to invest for preemption and the probability of innovating decrease accordingly.
Let us now consider how the mixed strategy equilibrium changes with the emission cap E.

Proposition 15. In the infinite delay case, the following equilibria are feasible:

- if $E < \varepsilon$, there are a symmetric mixed strategy equilibrium $((p^*, 1 - p^*), (p^*, 1 - p^*))$ and two asymmetric pure strategy Nash equilibria (*innovate*, *not innovate*), provided that it is also $\overline{C} < C < \widetilde{C}$;
- if $E > \varepsilon$, there is a pure strategy equilibrium (*not innovate*, *not innovate*).

First of all, changes in the emission cap affect the mixed strategy equilibrium because they affect the critical cost range. It can be $\overline{C} < C < \widetilde{C}$ only for $\theta < E < \varepsilon$. For $E > \varepsilon$,

---

[26] For each firm it is $V_N > V_I$ either if the other firm innovates for sure or if it does not. If one firm adopts k_1 with certainty, the other firm's expected payoff of not innovating collapses to Π_2^1, whereas its expected payoff of innovating degenerates to Π^2, with $\Pi_2^1 > \Pi^2$. If one firm does not adopt the new technology for sure, the other firm's expected payoff of not innovating is $\pi^0/(1-\rho)$, whereas its expected payoff of innovating degenerates to Π_1^1, with $\Pi_1^1 < \pi^0/(1-\rho)$ for $C > \widetilde{C}$.

the relative magnitude of per-period profits is such that the benefit of being first is always lower than the benefit of being second[27].

Therefore, the mixed strategy equilibrium arises when the switching cost and the emission cap are such that the payoff profiles imply a preemption game: each firm gets the highest payoff if it is the only one to innovate, and the lowest payoff if both do.

Proposition 16. In the infinite delay case, when a mixed strategy equilibrium arises, the probability of not innovating p^* is decreasing in the emission cap E, other things being equal.

This may appear counterintuitive. However, an increase in the permit supply, provided that it is not excessive (such that the condition $E < \varepsilon$ is met) makes the innovating firm's profit higher: the incentive to adopt first becomes stronger and then the probability of innovating goes up.

Appendix 2 reports a numerical example.

4 The Experiment

We designed and implemented an experiment that replicated the "innovation game" described in the previous chapter. The experiment was computerized using the Z-Tree software developed at the University of Zurich by Urs Fischbacher.

As in the "innovation game", the decision problem involves two subjects, representing the two duopolistic firms. Subjects play a dynamic game that ends after a random number of periods; we will explain later how subjects' decision problem in this setting is equivalent to the theoretical decision problem over an infinite horizon. In the initial round of the game, the two players are both in a state that we denote as "state A", which is associated with a symmetric combination of payoffs. Players have to decide whether to remain in "state A" or to switch to a state denoted as "state B". Once a player has decided to

[27] Recalling proposition 7:

- if $E < \varepsilon$, there exists an investment cost $C < \widetilde{C}$ such that the cost savings of waiting for one period are lower than the forgone extra-profit of being first, provided that the investment cost is sufficiently high to avoid imitation: a mixed strategy equilibrium is feasible;
- if $E > \varepsilon$ (the permit supply is sufficiently large), for every $C > \overline{C}$ the cost savings of delaying adoption of one period are higher than the extra-profit that can be realised by adopting technology k_1, and firms never innovate.

change state, he\she cannot return to the previous one. Each combination of states is associated with a combination of payoffs. Hence, “state A” corresponds to the old technology and “state B” corresponds to the new technology. The decision of moving from A to B corresponds to the decision of innovating. We leave the setting as abstract as possible, not using the words “technology” and “innovation”, in order to avoid any influence that this wording might have on subjects.

When deciding to change state, players incur in a once-for-all cost, which is deducted from their payoff. It is not explicitly stated that there is a switching cost; however this is evident from the payoff structure: players know their net payoffs associated to their decisions, so that they do not need to make any calculation to determine what they will be paid.

Subjects are paid the payoff corresponding to the combination of states they are in when the game finishes. The number of rounds is randomly determined and subjects do not know which round is going to be the final one. Each player’s decision problem ends when he\she has already moved from A to B, even though the game continues until the randomly determined number of rounds is over. The switching cost is incurred only once, in the round in which a subject decides to move from A to B: only if this round happens to be the final one, the subject that has moved will get a lower payoff (i.e. it will pay the once-for-all-cost), otherwise the subject will get the “full” payoff. In the instructions, all the payoffs are referred to as being *potential*, unless they are the payoffs corresponding to the state the subjects are in when the game ends (*actual* payoffs).

The instructions of the experiment are reported in Appendix 3. These instructions refer to the first treatment: instructions are the same for all the treatments, except for payoffs.

Let us now consider the equivalence between the experimental design, in which the game ends according to a random stopping rule, and the theoretical game, in which each firm’s life is infinite. In the theoretical analysis, firms discount their per period profits and the investment cost with a discount factor ρ in order to determine their lifetime profits. In the experiment, there is a probability λ that the game ends at each round, and a probability $(1-\lambda)$ that the game will go on to the next period[28]. Let Π_t denote the payoff that the subject would get if the game finished in period t. Then, as viewed from period 1

[28] Whichever the period players are in, they know that the game will finish after a finite number of repetitions but they do not know when: the number of rounds that remains to be played could be very large.

the probability of the game stopping in period 1 (and getting Π_1) is λ, the probability of it stopping in period 2 (and getting Π_2) is $\lambda(1-\lambda)$, the probability of it stopping in period t is $\lambda(1-\lambda)^{t-1}$. Hence, considering a stream of payoffs Π_1, Π_2,..., Π_t, each player expected payoff is

$$(\lambda)\sum_{t=1}^{\infty}(1-\lambda)^{t-1}\Pi_t \tag{9}$$

Therefore, a discount factor ρ in theory is equivalent to a probability $(1-\lambda)$ that at each round the game continues to the next one. The theoretical decision problem is to choose whether and when innovate in order to maximise the stream of future per period profits. In the experiment, a subject has to choose whether and when to change state in order to maximise (9). The two decision problems are equivalent; the only difference is the scaling factor λ, which does not affect decision.

In the experiment, the payoff functions and the probability λ are common knowledge. At the end of any one round each subject is told his\her rival's decision.

We implemented 4 treatments, differing one another in terms of the investment cost and\or the emission cap. Hence, C and E are our treatment variables. In particular, the values of C are such that two treatments imply quick imitation and the other two imply infinite delay, and the values of E are "high" for two treatments and "low" for the other two, as summarised in Table 1.

Table 1 - Treatments

C \ E	High $E = 22$	Low $E = 18.5$
Quick imitation $C = 1000$	T1	T2
Infinite delay $C = 1500$	T3	T4

Appendix 3 reports the payoff tables of the 4 treatments, that were given to subjects together with the instructions. The payoffs and the probability of the game stopping at any one round ($\lambda = 0.1$) are all that subjects needed to know in order to reach their decision.

Each treatment requires 9 subjects. 8 subjects, divided into 4 pairs, play the game 5 times, each of which with a different opponent. Hence, each subject faces the decision problem underlined above with 5 different players, without knowing who they are. By making the subjects playing several times we aim at collecting a sufficient number of observations that allows us to evaluate whether each individual always adopts the same strategy or randomises. The "absolute stranger" matching[29] should control for correlation between repetitions of the game: as far as strategic interaction is concerned, the strategy each subject chooses to play in a repetition should be independent on what happens in the previous one(s) since the subject pairs are different[30]. For the scope of the data analysis, we make the assumption that the 5 plays of the game in each treatment are independent one another.

The ninth subject does not play the game, but is in charge of determining the number of rounds each of the 5 games lasts. We call this subject the "Round Determinator" (Allsopp, 2002). The "Round Determinator" is elected by the subjects in each session and determines the number of rounds of each of the 5 repetitions of the game, as described in detail in the instructions. This number is not revealed to the other

[29] The "absolute stranger" matching implies that at any game each subject plays with a new partner, who is not known, and two subjects would never play together more than once.

[30] However, the 5 games will always be correlated, because each subject gains experience from one game to the next (learning).

participants in the experiment, though the subjects can check that the Round Determinator has performed the job described[31].

The experiment consisted of 4 sessions, one for each treatment. All the sessions took place in the Laboratory of the Centre for Experimental Economics (EXEC), at the University of York. Participants in the experiment were all undergraduate students, except for two postgraduate students. There were no trial periods.

Each game lasted 3-4 minutes depending on the number of rounds. At the end of the fifth play, the subjects were informed that the experiment was complete and were paid in cash and were then free to leave. Each session lasted between 25 and 45 minutes.

Table 2 reports for each treatment the number of times in which one subject was in state A whereas the other was in state B and the number of times the subjects in A followed immediately or waited some rounds to switch to state B. A time lag equal to 1 round means that the subject that was in state A moved to state B immediately the following round; a time lag equal to 2 means that the subject that was in state A waited 1 round before switching to B, and so on. The column labelled "never" reports the number of times that subjects who were in state A whereas their opponent was in state B were cut off in state A by the end of the game. We cannot be sure that if the game lasted longer they would have not moved.

Table 2 – Time lags before following

		Time lag						
Treatment	Subjects in A when the other is in B	1 round	2 rounds	3 rounds	4 rounds	5 rounds	7 rounds	never
1	7	5	0	0	0	0	0	2
2	13	6	5	1	0	0	1	0
3	13	7	2	1	0	1	0	2
4	7	5	1	0	0	0	0	1

No statistical analysis is needed to notice that subjects tended to quickly follow the player that had already moved in all the treatments: most observations are concentrated in time lags 1 and 2. There was a propensity to follow also in the infinite delay

[31] In this respect, this mechanism has the advantage of being transparent. Another advantage is that the number of rounds is determined by someone other than the experimenter, hopefully leading to a higher "trust degree" towards the mechanism itself. Moreover, this method is easy to carry out. However, there is always the risk that subjects who are not familiar with probabilities might not understand the mechanism.

treatments, despite the fact that the expected benefits from switching were 35, 4, -26 and -48 in treatments 1, 2, 3, and 4 respectively[32].

A possible explanation to this behaviour might be that there was a herd behaviour component: the players that had not moved to state B might have just imitated their rivals thinking that if the other players had moved that was the right thing to do.

Another possible explanation might be that subjects competed between each other: the players that had not moved were not willing to let their respective rivals earn more. Indeed, the risk of being caught by the end of the game was low and for a moderately risk loving person, this competitive (or "envy") component might affect behaviour[33].

All the 4 subject pairs were in state A at least 20 times in each treatment, since in the first round of each of the 5 plays all the subjects started in state A. In treatment 1 and 4 at least one subject in each subject pair moved to state B at the first round, whereas in the other two treatments some subject pair managed to coordinate on both not moving for more than 1 round. However, only in the 5th play of the second treatment a pair managed to coordinate on the (*not move to B*, *not move to B*) outcome, i.e. on the outcome with the highest expected payoff.

10 players over all the treatments always chose to move to state B when both subjects were in state A. We regard these subjects as playing a pure strategy "*move to B (innovate)*" in any round. Let us denote these players as "type B subjects". To be more precise, these subjects could also have adopted a mixed strategy, but in those plays it happened that they picked the B decision. We do not have enough observations to be reasonably sure that these players adopted a pure strategy[34]. On the other hand, there are no subjects that always chose to remain in state A, given that both players in the pair were in A. Hence, we conclude that no subjects played a pure strategy "*not move to B (not innovate)*".

Let us consider the probability of not moving to B over all the sample. Since we are interested in players' choice when both could make a decision, the total number of

---

[32] By switching state subjects could either gain an extra-profit with probability 0.9 or lose their payoff with a probability 0.1.

[33] If so, it might be interesting to run the experiment with a different parameter set, either increasing the probability that the game ends at each round, i.e. decreasing the discount factor, or reducing the difference between the payoff when both players are in state B and the payoff of the player that remains in state A. Another possibility is to do a treatment in which one of the players is the computer.

[34] It might be worthwhile to repeat the experiment and submit a questionnaire to subjects, asking them to describe their strategy, though this could change their behaviour

rounds in which both subjects were in state A represents the number of observations for each treatment, i.e. our sample sizes.

From each treatment, we estimate the probability that at least one player decides to remain in state A when both players are in A, that is Pr{*at least one remains in A* | *both in A*}. Denote these probabilities as $\hat{p}_1, \hat{p}_2, \hat{p}_3, \hat{p}_4$. Therefore, $\hat{p}_\tau$, $\tau = 1,\ldots,4$, is the sample proportion between the number of times at least one player's choice is A when both are in A, and the number of times both players are in A. We are interested in testing the hypothesis that each treatment sample probability is not statistically different from the corresponding probability suggested by the theory. Denote these latter probabilities as $p_1^*, p_2^*, p_3^*, p_4^*$. Given the sample sizes of our treatments, $n_1 = 40$, $n_2 = 54$, $n_3 = 54$, $n_4 = 40$ respectively, we can approximate the sample distribution of each proportion as

$$\hat{p}_\tau \sim N\left(p_\tau^*, \frac{p_\tau^*\left(1-p_\tau^*\right)}{n_\tau}\right)$$

The null hypotheses we want to test are $H_0 : \hat{p}_\tau = p_\tau^*$, $\tau = 1,\ldots,4$. The rejection rule for a two-tailed hypothesis test will lead us to reject H_0 when $|z_\tau| > c_\alpha$, where $z_\tau = \dfrac{\hat{p}_\tau - p_\tau^*}{\sqrt{\dfrac{p_\tau^*\left(1-p_\tau^*\right)}{n_\tau}}}$ is the usual z-score, c_α is the $(1-\alpha)\%$ critical value, and α is the significance level.

The results we get are summarised in the following table.

Table 3 - Estimates of the probability of not moving to B over all the sample

Treatment	n	$\hat{p}_\tau$	p^*	z	Result
1	40	0.25	0.84072	-10.2095	**
2	54	0.555556	0.66667	-1.73211	
3	54	0.611111	0.24572	6.236798	**
4	40	0.25	0.40506	-1.99767	*

In table 3, each row refers to a treatment. The columns show, for each treatment, the sample size, the sample probability, the theoretical probability, the z-score, and the results we get in terms of rejecting the null hypothesis, respectively. By "*" and "**" we indicate that we reject the null hypothesis that the two probabilities are not statistically different at 5%, and 1% significance level, respectively. If the cell reports no "star", then for these two significance levels we cannot reject the null hypothesis.

We estimate the probability of not moving to B excluding the type B subjects from the sample size of each treatment. As mentioned above, here we are making the assumption that the type B subjects played a pure strategy "move to B"; since no subject chose to stay in A in all the rounds in which both he\she and his\her opponent were in state A, there are no subjects who played a pure strategy "not move to B". Hence, the players that are not type B are assumed to have played a mixed strategy. Since we are interested in the choice of the subjects that played a mixed strategy when both could make a decision, the total number of rounds in which both subjects (of this type) were in state A represents the number of observations for each treatment, i.e. our sample sizes. Therefore, the analysis is analogous to the previous one, with different sample sizes of the 4 treatments, $n_1 = 25$, $n_2 = 44$, $n_3 = 49$, $n_4 = 20$.

For each treatment, let us denote the probability that at least one (mixed strategy) player decides to remain in state A when both (mixed strategy) players are in A as $\widehat{p}_\tau$, $\tau = 1,...,4$. As before, $\widehat{p}_\tau$, $\tau = 1,...,4$, is the sample proportion of the number of times at least one (mixed strategy) player's choice is A when both are in A. The null hypotheses we want to test are $H_0 : \widehat{p}_\tau = p^*_\tau$, $\tau = 1,...,4$. As we did for the other estimated probability $\hat{p}_\tau$, we apply the rejection rule for a two-tailed hypothesis test and we summarize the estimation results in the following table, which is analogous to table 3.

Table 4 - Estimates of the probability of not moving to B over the "mixed strategy" sample

Treatment	n	$\widehat{p}$	p^*	z	Result
1	25	0.4	0.84072	-6.0218	**
2	44	0.681818	0.66667	0.213152	
3	49	0.673469	0.24572	6.954958	**
4	20	0.5	0.40506	0.86494	

We can make the following considerations:

a) comparing the estimated probabilities $\widehat{p}_\tau$ and $\widehat{p}_\tau$, $\tau = 1,...,4$, we get the same results: we cannot reject the null hypothesis that the estimated probability is not statistically different from p^* only in treatments 2 and 4;

b) comparison between quick imitation and infinite delay treatments: for the same permit supply, the probability of not innovating should be lower in treatments 3 and 4, but the estimated probability follows this trend only for treatments 2 and 4, corresponding to the low permit supply;

c) comparison between high and low permit supply treatments: for the same cost, the probability of not innovating should be increasing in E in the quick imitation treatments and decreasing in E in the infinite delay ones, but the estimated probability does not fit this trend in any treatment.

We conclude this section, observing that the most frequent final outcome was (*State B*, *State B*) in all the treatments. Subjects changed state more often then predicted by the theory, except than in treatment 3, and they generally followed if their respective opponents had switched to state B. Players appeared to be "attracted" by changing state for reasons that are out of the strategic structure. Both players in state B (in any period) is the Pareto inferior pure strategy equilibrium in the quick imitation case; hence, we conclude by saying that subjects failed in coordinating on the Pareto superior equilibrium in each period. On the other hand, both players in state B is not a pure strategy equilibrium in the infinite delay case, since in this case pure strategy equilibria are asymmetric. However, (B,B) is the most likely outcome when subjects play a symmetric mixed strategy, where the probability of (B,B) arising in any period is $(1-p^*)(1-p^*)$, equal to 0.569 and 0.354 in treatment 3 and 4, respectively.

Concluding remarks

We have looked at the impact that a market of emission permits may have on the propensity of firms to invest in environmental friendly technologies. In particular, we have focused on the effects of the interaction between the output and the permit markets on firms' investment decisions. We have looked at the conditions under which the diffusions of a cleaner technology in a non-competitive setting is more likely to occur. In this respect, we aimed at deriving some suggestions for environmental regulators. We have addressed this problem from both a theoretical and an experimental point of view.

In a non-competitive setting, firms' investment decisions have a strategic component, since a firm can use the adoption of a cleaner technology to enlarge its market share in the industry at the expense of its rival. We modelled this strategic interaction as a two-firm dynamic game, in which each firm has to choose whether and when to make an irreversible investment, that affects the infinite stream of its future profits. We solved this "innovation game" looking for symmetric stationary equilibria. We have seen that the stationary equilibria to this game crucially depend on both the cost of switching to the cleanest technology and the emission cap. Given the directions in which the investment cost and the emission cap affect the feasible stationary equilibria, an environmental regulator aiming at speeding up the diffusion of an environmental friendly technology has to adjust these variables in the appropriate manner.

It is intuitive that, given the emission cap, subsidising the duopolists in order to lower the cost of changing technology will induce both firms to adopt the superior technology. However, we can expect this "diffusion" outcome only if the cost of switching is sufficiently low. In particular, the cost must be such that the cost savings from delaying adoption are lower than the net benefits from being the leader.

For some higher values of the investment cost, this "diffusion" outcome is only one of the possible outcomes. However, the control authority can still use the investment cost and the permit supply as instruments to push firms towards the most desirable outcome, from an environmental point of view. In particular, given the permits supply, the regulator should try to push the investment cost as close as possible to the critical value that splits the quick imitation and the infinite delay cases; this is the cost below which the savings from delaying imitation are lower than the gains from following and above which the opposite occurs. The closer is the switching cost to this value, the higher is the probability that any one firm innovates in a period. If the investment cost is above this critical value, the regulator can make joint adoption more likely by subsidising firms. Alternatively, given the cost, it can increase the permit supply such that the benefit from switching is higher than the benefit from delaying adoption. In this case the objective is to make innovation as profitable as possible, increasing both firms preemption incentive. On the contrary, if the investment cost is below this critical value, the regulator may expect a quick imitation and can push diffusion by taxing the investment! Alternatively, given the cost, it can decrease the permit supply. In this case the objective is to make firms afraid of being preempted, so that each firm is not induced to adopt because it will be better off by doing so, but it will be tempted to invest in order to avoid to be worse off if its rival

moves first. In other words, the regulator objective is to make more difficult for firms to coordinate on the non adoption outcome (without communication).
For the cost and emission cap values for which the “diffusion” equilibrium is only one of the possible stationary equilibria, the diffusion is not the dominant outcome: firms could be better off by doing something else, either coordinating on non adoption or on an outcome in which only one adopts.
Game theory does not help us in predicting which outcome will actually occur. In particular, not necessarily firms will manage to coordinate on the Pareto superior one. An experimental investigation can help us in seeing this.
Our pilot experiment provided interesting initial results as well as revealing some points in the experimental design and parameter set that could be worthy to revise. The mostly observed equilibrium was not the Pareto dominant one. This result is consistent with other experimental investigations, which support the conclusion that a Pareto dominant equilibrium does not necessarily represent a focal one (Ochs, 1995).
In interpreting these initial results we must pay special attention; in particular, we should consider the following:

- subjects might not have completely understood the round determination mechanism or they might not have believed it, so that this mechanism failed in inducing discounting: subjects looked at payoffs associated to each combination of states, rather than at the expected payoffs;
- subjects might well understood the random number of rounds in each play, but the probability of stopping the play at each round was too low, so that subjects did not consider the risk associated with changing state;
- the difference between payoffs corresponding to different outcomes is not striking, even though parameters lead to substantially different p^*: probably, the experiment should be designed in such a way that it penalizes bad play sufficiently, especially as far the imitation choice is concerned;
- it is difficult to implement a mixed strategy.

Further experimental sessions are needed in order to both provide a larger number of observation for statistical analysis and test subjects’ behaviour under different parameter sets. If these results prove to be robust to changes in parameter set and further experimental sessions, we could conclude that (*innovate*, *innovate*) is the most likely outcome of this “innovation game” when the permit supply and the investment cost are such that a stationary symmetric mixed strategy equilibrium may arise in any period.

This is an encouraging result from the environmental regulator’s point of view. When, given the permit supply, the investment cost is sufficiently low that each firm can gain by adopting first, even if it anticipates that its rival would follow, the environmental regulator can “predict” that both firms will innovate straight away. If the investment cost is even slightly higher than this critical value, the outcome of the game will be “unpredictable”. The experiment focused on these “unpredictable” cases, showing that even in these cases the control authority does not need to worry too much, since firms’ behaviour will eventually lead to diffusion. The experiment results suggest that when quick imitation is expected, firms may fail to coordinate on not innovating, since for each of them the fear of the other innovating first apparently prevails. In the infinite delay case, when a preemptive equilibrium is the dominant one, the preempted firm does not apparently leave the other firm maintain this advantage position, and tend to imitate, even though for a firm (maximising its expected payoff) would be better not to do so.

References

Allsopp, L. (2002). “Search and externalities. A pilot experiment”. Working Paper, Adelaide University .

Baumol, W. J. and Oates W. E. (1988). *The Theory of Environmental Policy*, second edition, Cambridge University Press.

Denicolo, V. (1999). “Pollution-Reducing Innovations under Taxes or Permits”, *Oxford Economic Papers*, vol. 51(1), pp. 184-199.

Downing, P. B. and White, L. J. (1986). “Innovation in pollution control”, *Journal of Environmental Economics and Management*, vol. 13, pp. 18-29.

Fershtman, C. and de-Zeeuw, A. (1995), “Tradeable Emission Permits in Oligopoly”, Tilburg CentER for Economic Research Discussion Paper.

Fischer, C., Parry, I. W. H. and Pizer, W. A. (1999) “Instrument choice for environmental protection when technological innovation is endogenous”, Resources for the Future Discussion Paper 99-04.

Fudenberg, D. and Tirole, J. (1985). “Preemption and rent equalization in the adoption of new technology”, *Review of Economic Studies*, vol. 52, pp. 383-401

Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press, Cambridge, Massachusetts.

Jaffe, A. B. and Stavins, R. N. (1995). “Dynamic incentives of environmental regulations: the effects of alternative policy instruments on technology diffusion”. *Journal of Environmental Economics and Management*, vol. 29 (3suppl. Part 2), pp. S43-S63.

Jaffe A.B., Newell R.G., Stavins R.N. (June 2002), “ Environmental Policy and Technological Change”, *Environmental and Resource Economics*, vol. 22, pp. 41-70.

Jung, C., Krutilla, K. and Boyd, R. (1996). “Incentives for advanced pollution abatement technology at the industry level: an evaluation of policy alternatives”, *Journal of Environmental Economics and Management*, vol. 30, pp. 95-111.

Kagel, J. H. and Roth, A. E. (1995, eds.). *The Handbook of Experimental Economics*, Princeton University Press, Princeton, New Jersey.

Laffont, J. J. and Tirole, J. (1996a). “Pollution permits and compliance strategies”, *Journal of Public Economics*, vol. 62, pp. 85-125.

Laffont, J. J. and Tirole, J. (1996b). “Pollution permits and environmental innovation”, *Journal of Public Economics*, vol. 62 (1-2), pp 127-140.

Magat, W. A. (1978). “Pollution control and technological advance: A model of the firm”, *Journal of Environmental Economics and Management*, vol. 15, pp. 1-25.

Maleug, D. A. (1989). “Emission trading and the incentive to adopt new pollution abatement technology”, *Journal of Environmental Economics and Management*, vol. 16, pp. 52-57.

Milliman, S. R. and Prince, R. (1989). “Firm incentives to promote technological change in pollution control”, *Journal of Environmental Economics and Management*, vol. 17, pp. 247-265.

Montero J-P. (2002), “Permits, Standards, and Technology Innovation”, *Journal of Environmental Economics and Management*, vol. 44 (1), pp. 23-44.

Montgomery, W. D. (1972). ‘Markets in licenses and efficient pollution control programs’, *Journal of Economic Theory*, vol. 5, pp. 395-418.

Ochs, J. (1995). “Coordination problems”, in Kagel J.H. and Roth, A. E. (1995,eds.), 195-252

Requate, Till. (1993). “Pollution Control in a Cournot Duopoly via Taxes or Permits”, *Journal of Economics*, vol. 58(3), pp. 255-291.

Requate, T. (1998). “Incentives to Innovate under Emission Taxes and Tradeable Permits”, *European Journal of Political Economy*, vol. 14(1), pp. 139-65.

Reinganum, J. F. (1981). “On the diffusion of new technology: a game theoretic approach”, *Review of Economic Studies*, vol. 48, pp. 395-405

Tietenberg, T. H. (1985). *Emission Trading. An Exercise in Reforming Pollution Policy*. Resources For The Future, Washington, D.C.

Tirole, J. (1988). *The Theory of Industrial Organization*. MIT Press, Cambridge, Massachusetts

Appendix 1 - Figures

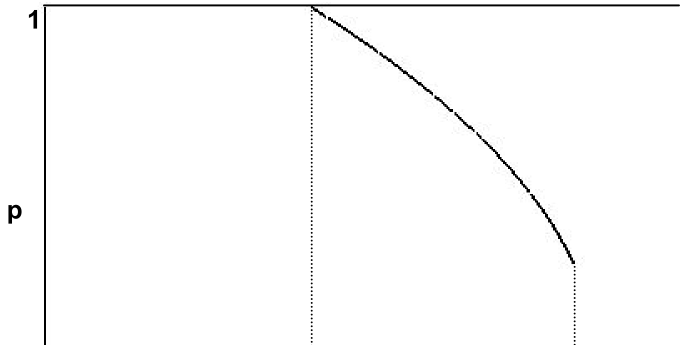

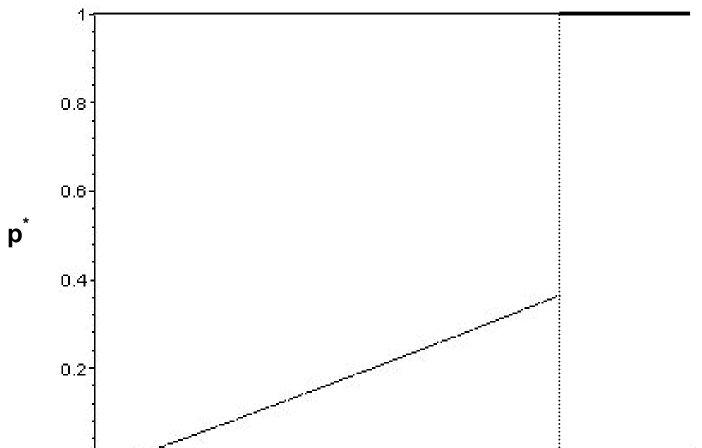

Figure A1.2 – Infinite delay: optimal probability of not innovating in terms of the investment cost, $p^* = p^*(C)$, for fixed emission cap and technology parameters.

Appendix 2 – Numerical Examples

Table A2.3 reports a numerical example for the quick imitation case, based on the parameters reported in Table A2.1. Table A2.2 reports the equilibrium permit prices, the per-period profits and the critical cost values implied by this parameter set. The equilibrium probability and the payoffs are calculated for increasing values of the investment cost C and for two different values of the permit supply E. The values of C are the same for the two emission caps.

Table A2.1 - Parameters

a	b	c	k_1	k_2	ρ	θ	α^2	ε
10	1	2	0.7	0.8	0.9	0.6222	3.733	2.987

Table A2.2 - Permit price, per period profits and critical cost for two values of the permit supply

E	q^0	q^1	q^2	π^0	π_1^1	π_2^1	π^2	$\overline{C}$	$\hat{C}$	$\tilde{C}$
2	5.3125	5.2632	5.3061	1.5625	2.6051	1.1831	2.0408	8.5768	5.347	10.4261
2.5	4.1406	3.9474	3.7755	2.4414	3.5239	2.1977	3.1888	9.9110	7.808	10.8245

The following can be noticed:

- when the permit supply is higher, the critical values for the cost, $\hat{C}$ and $\overline{C}$, increase; therefore for $E = 2.5$ only the highest values of C leads to a solution between 0 and 1;
- the probability of not innovating p^* is higher than in the infinite delay case;
- p^* is decreasing in C and increasing in E;
- lifetime profits and expected payoffs decrease with C and increase with E;
- the expected payoff implied by p^* is increasing in p^*, and so it is decreasing in C;
- the expected payoff of never innovating is constant;
- for all the values of C and E, it is $\Pi_1^1 > \Pi^2 > \Pi_2^1$ and $\frac{\pi^0}{1-\rho} > V(p^*) > \Pi^2$

Table A2.4 reports a numerical example for the infinite delay case, based on the parameters reported in Table A2.1. The equilibrium probability and the payoffs are calculated for increasing values of the investment cost C and for two different values of the permit supply E. The values of C are the same for the two emission caps. The following can be noticed:

- when the permit supply is higher, the critical values for the cost, $\overline{C}$ and $\widetilde{C}$, increase; therefore for $E = 2.5$ only the highest values of C lead to a solution between 0 and 1 (for lower C, the game is a quick imitation one);
- the probability of not innovating p^* is lower than in the quick imitation case;
- p^* is increasing in C and decreasing in E;
- lifetime profits decrease with C and increase with E;
- the expected payoff of never innovating is constant;
- the expected payoff implied by p^* is increasing in p^*, and so it is increasing in C.
- for all the values of C and E it is $\Pi_1^1 > \Pi_2^1 > \Pi^2$ and $\Pi_1^1 > V(p^*) > \Pi_2^1$.

Table A2.3 - Quick imitation case: numerical example

C	$E = 2$						$E = 2.5$					
	p^*	Π_1^1	Π_2^1	Π^2	$\frac{\pi^0}{1-\rho}$	$V(p^*)$	p^*	Π_1^1	Π_2^1	Π^2	$\frac{\pi^0}{1-\rho}$	$V(p^*)$
5.5974562	0.9693337	15.375	14.51277	14.810707	15.625	15.357695	-	26.625377	25.858945	26.290299	24.414063	-
5.8474562	0.9374234	15.125	14.28777	14.560707	15.625	15.089688	-	26.375377	25.633945	26.040299	24.414063	-
6.0974562	0.9041042	14.875	14.06277	14.310707	15.625	14.820887	-	26.125377	25.408945	25.790299	24.414063	-
6.3474562	0.8691712	14.625	13.83777	14.060707	15.625	14.551174	-	25.875377	25.183945	25.540299	24.414063	-
6.5974562	0.8323644	14.375	13.61277	13.810707	15.625	14.280404	-	25.625377	24.958945	25.290299	24.414063	-
6.8474562	0.7933453	14.125	13.38777	13.560707	15.625	14.008386	-	25.375377	24.733945	25.040299	24.414063	-
7.0974562	0.7516596	13.875	13.16277	13.310707	15.625	13.734863	-	25.125377	24.508945	24.790299	24.414063	-
7.3474562	0.7066719	13.625	12.93777	13.060707	15.625	13.459477	-	24.875377	24.283945	24.540299	24.414063	-
7.5974562	0.6574457	13.375	12.71277	12.810707	15.625	13.181699	-	24.625377	24.058945	24.290299	24.414063	-
7.8474562	0.6024948	13.125	12.48777	12.560707	15.625	12.900691	0.9924072	24.375377	23.833945	24.040299	24.414063	24.372833
8.0974562	0.5391864	12.875	12.26277	12.310707	15.625	12.614966	0.9415823	24.125377	23.608945	23.790299	24.414063	24.105802
8.3474562	0.4619046	12.625	12.03777	12.060707	15.625	12.321357	0.8872552	23.875377	23.383945	23.540299	24.414063	23.837599

Table A2.4 - Infinite delay case: numerical example

C	$E = 2$						$E = 2.5$					
	p^*	Π_1^1	Π_2^1	Π^2	$\frac{\pi^0}{1-\rho}$	$V(p^*)$	p^*	Π_1^1	Π_2^1	Π^2	$\frac{\pi^0}{1-\rho}$	$V(p^*)$
8.8268306	0.0476502	17.224262	11.831333	11.581333	15.625	11.85022	-	26.411704	21.976762	23.060925	24.414063	-
9.0768306	0.0956419	16.974262	11.831333	11.331333	15.625	11.871033	-	26.161704	21.976762	22.810925	24.414063	-
9.3268306	0.1440352	16.724262	11.831333	11.081333	15.625	11.894113	-	25.911704	21.976762	22.560925	24.414063	-
9.5768306	0.1929064	16.474262	11.831333	10.831333	15.625	11.91989	-	25.661704	21.976762	22.310925	24.414063	-
9.8268306	0.2423546	16.224262	11.831333	10.581333	15.625	11.948923	-	25.411704	21.976762	22.060925	24.414063	-
10.076831	0.2925115	15.974262	11.831333	10.331333	15.625	11.981955	0.0535876	25.161704	21.976762	21.810925	24.414063	21.990485
10.326831	0.3435578	15.724262	11.831333	10.081333	15.625	12.020005	0.1353086	24.911704	21.976762	21.560925	24.414063	22.014314

Appendix 3 – Instructions and payoff matrices

INSTRUCTIONS OF THE EXPERIMENT

Welcome to this experiment. The instructions are simple. If you follow them carefully, you could make a considerable amount of money, which will be paid to you in cash immediately after the experiment. You should note that you should not talk to the other participants during the experiment. If you do, it will defeat the purpose of the experiment, and we will have to ask you to leave.

The experiment consists of 5 plays of a game, each of which involves you and some other participant in the experiment. This other person will be changed between each play of the game, and you will never know who he or she is, and he or she will not know who you are. What you do in the experiment will not be divulged to anyone.

The game is simple. It will last a randomly determined number of rounds. We will explain first how this number of rounds is determined, and then we will explain the structure of the game that you will be playing.

The participants in the experiment will elect one of their number as the *Round Determinator*. This participant will be paid a fixed fee of £10 for participating in the experiment. The Round Determinator will not play the game but will determine the number of rounds in each game. He or she will do this as follows. Before each game the Round Determinator will go with one of the experimenters into the side office where there is an opaque bag containing 9 *blue* balls and 1 *white* ball. The Round Determinator will shake the bag and pick one ball from it at random, note the colour, and then replace the ball in the bag. He or she will do this repeatedly – until the *white* ball is drawn. The number of draws that is required is the number of rounds that the particular game will last. After the game, the Round Determinator will certify to the other participants that this procedure was followed, though obviously the Round Determinator will not confirm the number of rounds until after the game is finished. This procedure will be repeated for each of the 5 games. You should note that this procedure implies that, at the end of any one round, there is a 1 in 10 chance that the game will finish after that round and a 9 in 10 chance that the game will continue into the next round. Note also that the number of rounds in a game will vary randomly from game to game.

We now describe the game that you will be playing. You, and the person with whom you are playing, start the game in a particular state, which we call State A. In any round of the game, you can choose to change to a new state, which we call State B. Once you have changed, you cannot change back. So the decision-problem is simple: **all you have to decide is whether, and when, you want to change from State A to State B.** The other person with whom you are playing has exactly the same decision problem. In any round, you decide simultaneously, without knowledge of what the other player is doing. Please note that whichever the state you are in, the game will go on until it reaches its randomly determined end.

We now describe how you will be paid for participating in this experiment. You will receive a payment for each of the 5 games, and your payment for the experiment as a whole will be the sum of the payments for the 5 games. This final amount will be added to your participation fee of £2.

In any one game the payment is determined in the following way. In each round of a game, you have a *potential* payoff. This depends on the state that you and the other player are in that round, and also whether you have decided to change state that round. If you have *not* changed state in a particular round, then the following table gives your potential payoff in that round.

Your *potential* payoff if you have *not* changed state in this round		**If the other player is in State A**	**If the other player is in State B**
	If you *are* in State A	**1398**	**1315**
	If you *are* in State B	**1670**	**1500**

If you *have* changed state in a particular round, then the following table gives your potential payoff in that round.

Your *potential* payoff if you *have* changed state in this round		**If the other player is in State A**	**If the other player is in State B**
	(You *are* in State B)	**170**	**0**

These potential payoffs are measured in *tokens*. The tokens you will earn during each game will be converted in cash at the end of the experiment at the conversion rate of *1000 tokens for £1*.

At the end of each round you will be told in which state you and the other player are and what your and the other player's *potential* payoffs for that round are.

Your *actual* payoff in a particular game is simply given by your potential payoff ***in the round in which the game finishes.***

We repeat that your actual payoff for any game will be the potential payoff in the round in which the game finishes. If you have not changed state in that final round, the first table above gives your actual payoffs; if you have changed state in that final round, the second table above gives your actual payoffs.

On your desk you have a page titled "**The Payoff Tables**" reporting the two tables above.

Your payment for the experiment as a whole will be the sum of your actual payoffs in all 5 plays of the game, converted into pounds (at the exchange rate of 1000 tokens = £1), plus the participation fee.

If you are unclear about any aspect of these instructions, please raise your hand and one of the experimenters will answer your question.

THE PAYOFF TABLES (treatment 1)

TABLE 1

		If the other player is in State A	If the other player is in State B
Your *potential* payoff if you have *not* changed state in this round	If you *are* in State A	1398	1315
	If you *are* in State B	1670	1500

TABLE 2

		If the other player is in State A	If the other player is in State B
Your *potential* payoff if you *have* changed state in this round	(You *are* in State B)	170	0

THE PAYOFF TABLES (treatment 2)

TABLE 1

		If the other player is in State A	If the other player is in State B
Your *potential* payoff if you have *not* changed state in this round	If you *are* in State A	1428	1346
	If you *are* in State B	1690	1500

TABLE 2

		If the other player is in State A	If the other player is in State B
Your *potential* payoff if you *have* changed state in this round	(You are in State B)	190	0

THE PAYOFF TABLES (treatment 3)

TABLE 1

		If the other player is in State A	If the other player is in State B
Your *potential* payoff if you have *not* changed state in this round	If you *are* in State A	1432	1376
	If you *are* in State B	1613	1500

TABLE 2

		If the other player is in State A	If the other player is in State B
Your *potential* payoff if you *have* changed state in this round	(You *are* in State B)	113	0

THE PAYOFF TABLES (treatment 4)

TABLE 1

		If the other player is in State A	If the other player is in State B
Your *potential* payoff if you have *not* changed state in this round	If you *are* in State A	1452	1398
	If you *are* in State B	1627	1500

TABLE 2

		If the other player is in State A	If the other player is in State B
Your *potential* payoff if you *have* changed state in this round	(You *are* in State B)	127	0