MANAGEMENT SCIENCE
Vol. 22, No. 7, March, 1976
Printed in U.S.A.

ON THE CONCEPT OF THE VALUE OF INFORMATION IN COMPETITIVE SITUATIONS*†

JEAN-PIERRE PONSSARD‡

Ecole Polytechnique

In this paper a game theoretic model is used to extend information value theory, as developed in decision analysis, to competitive situations. One of the main differences between competitive and noncompetitive situations is that part of the environment (namely the competitors) may be modified as a result of experimentation in another part of the environment (nature). Hence, states of the world and actions may no more be independent. Nevertheless, we shall show how the classical concept may be generalized to cover competitive situations.

1. Introduction

The concept of the value of information is one of the cornerstones of decision analysis [1], [2], [7]. It is intended to be a guide for the research and development of new strategies; in particular, for strategies which would allow for the gathering of new information on the real state of nature. However, in competitive situations such strategies may induce a change in the behavior of the competitors if these become aware of the experimentation. Then information usage is likely to become more complicated since a strategy used by the informed competitor may be used as a "second stage" experiment on nature by the uninformed competitors. So, for the decision maker who is interested in the value of an experiment which implies a modification of his behavior as perceived by competitors, strategies and states of nature (which include the competitors' strategies), may no longer be considered as independent as usually assumed in decision theory.

The present paper analyzes competitive situations in which individual experimentation is performed *with full knowledge of the competitors though the outcome is known only to the experimenter*. The analysis is based on a general game theoretic model developed by Harsanyi [1]. Since this paper is rather conceptual, it deals mainly with interpretations and discussions, relying on other papers for basic mathematical proofs [4], [6].

In §2 we shall define the concept of the value of information as used in decision analysis. In §3 we shall show how the concept may be extended to strictly competitive situations. This will be illustrated by means of an example in §4. In the last section some further remarks will be discussed including a nonstrictly competitive situation in which the value of information is negative.

2. The Value of Information Revisited

Consider the following classical decision problem under uncertainty: select an action among a finite set of feasible actions $A = \{a\}$, given a finite set of possible events or states of nature, $E = \{e\}$, a probability distribution on the events p_0

* Processed by Professor Ronald A. Howard, Departmental Editor for Decision Analysis; received October, 26, 1973; revised June 26, 1975. This paper has been with the author 2 months for revision.

† Part of this work was done while the author was with the International Institute of Applied Systems Analysis, Vienna, and appeared as "The Value of Information in Strictly Competitive Situations," Research Memorandum 74-11, *IIASA* (June 1974).

The author would like to express his thanks to John Harsanyi, Howard Raiffa and Robert Winkler for many interesting discussions on the subject.

‡ Chargé de Recherche au Groupe de Recherche en Gestion des Organisations.

$= \{p_0^e\}_{e \in E}$ ($p_0^e > 0$; $\sum_{e \in E} p_0^e = 1$), and a payoff function u (or more generally a utility function) defined on $A \times E$. (A and E are assumed finite for mathematical simplicity.)

According to decision theory, the selected action should maximize the expected payoff. Taking the probability distribution on E as a parameter $p \in P = \{p^e \mid p^e \geqslant 0, \sum_{e \in E} p^e = 1\}$, the optimal expected payoff $\bar{u}(p)$ is then obtained as

$$\bar{u}(p) = \operatorname*{Max}_{a \in A} \sum_{e \in E} u(a, e) p^e. \tag{2.1}$$

Let an experiment I^0 be defined as a random variable on P. Specifically assume that this random variable may take only a finite set of values $\{p_i\}_{i \in I}$ in P with respective probabilities γ_i($\gamma_i > 0$; $\sum_{i \in I} \gamma_i = 1$). For consistency we have

$$\sum_{i \in I} \gamma_i p_i = p_0. \tag{2.2}$$

An experiment may equivalently be defined by a matrix $Q = \{q_{ei}\}_{e \in E, i \in I}$ in which $q_{ei} = \text{Prob}\{i \mid e\}$. One may go from one definition to the other one by means of Bayes' theorem. We shall mostly use the first definition.

The expected value of the information to be revealed by the experiment I^0, $\text{EVI}(p_0 \mid I^0)$, is then defined as the incremental gain obtained by making one's decision depend on the outcome of the experiment. Namely

$$\text{EVI}(p_0 \mid I^0) = \sum_{i \in I} \gamma_i \bar{u}(p_i) - \bar{u}(p_0). \tag{2.3}$$

As a special case the expected value of perfect information, $\text{EVPI}(p_0)$, obtained by the experiment which would reveal the state of nature, is such that

$$\text{EVPI}(p_0) = \left[\sum_{e \in E} p_0^e \operatorname*{Max}_{a \in A} u(a, e) \right] - \bar{u}(p_0). \tag{2.4}$$

The expected value of information is generally interpreted as the maximal amount at which one would be willing to buy the experiment.

In the remainder of this section we shall state a simple property suggested by (2.3), (see proposition 2.1.1 in [5] for proof as well as some further development).

Denote by P_I the smallest convex subset of P which contains the vectors $\{p_i\}_{i \in I}$ and by $\text{Cav}_{P_I} f(p)$ the minimal concave function[1] greater or equal to $f(p)$ on P_I, in which $f(p)$ is any real-valued continuous function on P. Let $\text{Cav}_{P_I} f(p)|_{p_0}$ stand for the value of the function $\text{Cav}_{P_I} f(p)$ at p_0.

PROPOSITION 2.1. *For any experiment I^0 and any $p_0 \in P_I$, the value of information, $\text{EVI}(p_0 \mid I^0)$, is such that*:

$$\text{EVI}(p_0 \mid I^0) \leqslant \operatorname*{Cav}_{P_I} \bar{u}(p)\Big|_{p_0} - \bar{u}(p_0), \tag{2.5}$$

and the equality holds for perfect information, that is:

$$\text{EVPI}(p_0) = \operatorname*{Cav}_{P} \bar{u}(p)\Big|_{p_0} - u(p_0). \tag{2.6}$$

This result has a simple geometric interpretation. Indeed, assume that $E = \{e_1, e_2\}$ and let $\bar{u}(\cdot)$ be the optimal payoff function on $P = \{p = (p^1, p^2) \mid p^1 \geqslant 0, p^2 \geqslant 0, p^1 + p^2 = 1\}$ ($\bar{u}(\cdot)$, a convex function, is piece wise linear since the set of actions A is

[1] $g(p)$ is a concave function on P if and only if for all p_1 and p_2 and all $\lambda \in (0, 1)$: $g(\lambda p_1 + (1 - \lambda) p_2) \geqslant \lambda g(p_1) + (1 - \lambda) g(p_2)$.

finite). Let $p_0 \in P$ be the a priori probability distribution on E and the experiment $I^0 = \{i, j\}$ be defined by two possible a posteriori probability distributions on E, $p_i \in P$ and $p_j \in P$, with marginal probabilities γ_i and γ_j respectively (recall that for consistency we have $\gamma_i p_i + \gamma_j p_j = p_0$). Then the information value analysis is completely described by the following graph.

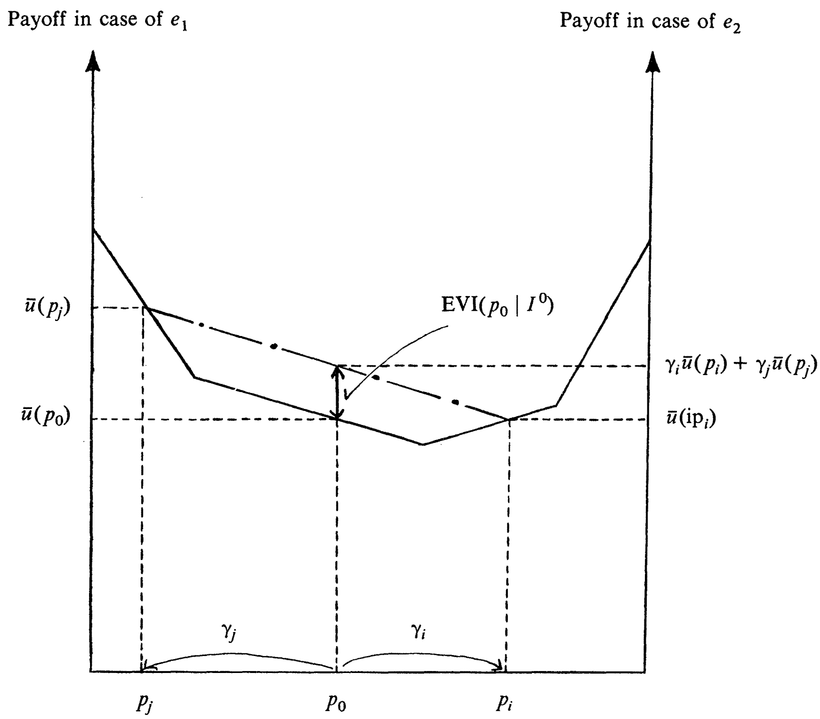

FIGURE 1. Geometric Interpretation of the Value of Information.

Since the function $\bar{u}(\cdot)$ is convex it may appear that the technical apparatus developed so far is unduly complicated. However, as we shall now show, it will turn out to be particularily well suited for the study of competitive situations.

3. Sequential Strictly Competitive Situations

In this section, we shall generalize the concept of the value of information, as recalled in §2, to the simplest form of competition; that is, the *constant sum* case.

Let the two competitors be competitor 1 and competitor 2, 1 selecting an action from A, and 2 from $B = \{b\}$. For any event, $e \in E$, we assume that the two competitors' payoffs, which are now defined on $A \times B \times E$, add up to some constant $c(e)$, independently of the selected actions. We assume that the two competitors move *sequentially*, 1 moving first; that is, 1 selects some action a which is revealed to 2, and then 2 selects some action b, both decision makers being uncertain about the event e which will prevail but having the *same probability distribution on E*. Then 1 gets $u(a, b, e)$ and 2 gets $v(a, b, e)$ such that $(a \in A)\,(b \in B)\; u(a, b, e) + v(a, b, e) = c(e)$.

Notice that, although the $c(e)$'s may be different so that the game in extensive form is nonconstant, the resulting game in normal form is constant sum. Namely we have, in terms of expected payoff

$$(a \in A)\,(b \in B) \qquad \sum_{e \in E} \{u(a, b, e) + v(a, b, e)\}\, p_0^e = \sum_{e \in E} c(e) p_0^e.$$

In these conditions, $\bar{u}(p_0)$, 1's optimal expected payoff is

$$\bar{u}(p_0) = \operatorname*{Max}_{a \in A} \operatorname*{Min}_{b \in B} \sum_{e \in E} p_0^e u(a, b, e), \tag{3.1}$$

and 2's optimal expected payoff is

$$\bar{v}(p_0) = \operatorname*{Min}_{a \in A} \operatorname*{Max}_{b \in B} \sum_{e \in E} p_0^e v(a, b, e). \tag{3.2}$$

These optimal payoffs are derived under the usual assumption that both competitors behave rationally so that competitor 2 maximizes his payoff conditional on the action selected by competitor 1, and competitor 1 selects his own action accordingly.

In this framework, what is the value of perfect information on E to competitor 1, *assuming that the other one will know*[2] that perfect information has been bought? Competitor 2, by observing competitor 1's selected action, may learn something about the state of nature observed by 1. How does this learning procedure operate, and what are its implications for information usage? These are the problems we now wish to investigate.

This investigation relies on a theoretical result proved in [6] (Theorem 1, page 101). In the context of this paper the result appears as an extension of Proposition 2.1.

PROPOSITION 3.1. *In a strictly competitive sequential situation the value of perfect information to competitor* 1 *may be expressed as*

$$\text{EVPI}(p_0) = \operatorname*{Cav}_{P} \bar{u}(p)\Big|_{p_0} - \bar{u}(p_0). \tag{3.3}$$

Insights provided by this result and their interpretation will be conveyed by means of an example. Let us however note immediately that in spite of the formal parallelism between (2.6) and (3.3), a significant difference lies in the fact that in (3.3) $\bar{u}(p)$ need not be convex. The implications of this fact for information usage will clearly appear in the example.

4. An Example

4.1. The Case.

Suppose that 1 and 2, the two competitors, have to set a price, $a \in A$ for 1 and $b \in B$ for 2, for a new product. Moreover, suppose that the size of the market, $e \in E$, is uncertain. Suppose also that 1 is the price leader so that 2 will wait until 1 has set up his price.

Assume that the payoff tables appear as follows:

(i) In case of a bad market, the benefits would add up to 6 and, depending on the prices set, would be shared such that:

Bad Market $e = e_1$		2's price: low	2's price: high
1's price	low	(5, 1)*	(1, 5)
	high	(3, 3)	(2, 4)

* (1's payoff, 2's payoff)

[2] There are certainly other ways to define the value of information in competitive situations. This specific assumption is particularly well suited in the context of game theory since it keeps the situation as a game with complete information; that is, both players know how much is known by the other one.

(ii) In case of a good market, the figures would add up to 9 and be such that:

Good Market $e = e_2$		2's price low	2's price high
1's price	low	(5, 4)	(6, 3)
	high	(4, 5)	(7, 2)

If there were no uncertainties, then the two competitors would sequentially set a high price (H) in case of a bad market and a low price (L) in case of a good market.

If they are uncertain about the market, then the prices to be set will depend on the probability distribution over E. These optimal prices and the associated payoff to 1 are depicted on Figure 2.

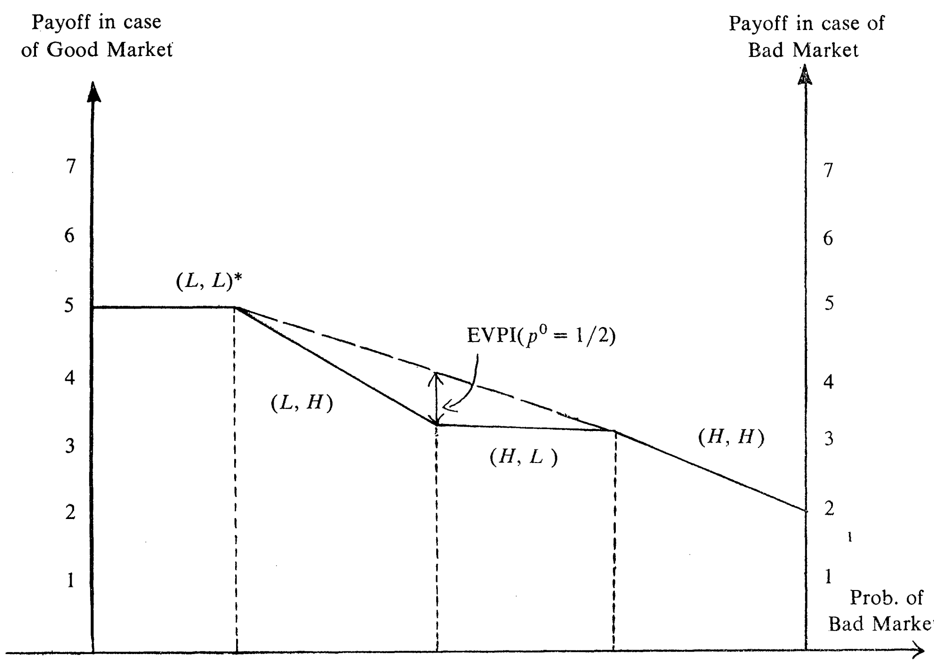

FIGURE 2. 1's optimal expected payoff.

One can see that if the probability of a bad market is less than $1/2$ competitor 1 should set a low price, and if it is greater than $1/2$ he should set a high price. Competitor 2 would follow competitor 1's price if the probability of a bad market is less than $1/5$ or greater than $3/4$. Between these two values competitor 2 would set the opposite price of competitor 1. Intuitively if the uncertainties are high, competitor 2 has much more to gain by taking a bold risk than by being a follower (for instance, suppose 1 sets a low price in the expectation of a good market, by setting a high price 2 may loose one unit if 1's expectation turns out to be wrong). We shall concentrate our analysis in the case of high uncertainties ($1/5 \leqslant p \leqslant 3/4$).

Suppose now that competitor 1 may order a market study and thus obtain perfect information while competitor 2 would only know the potential outcomes of the market study and remain uncertain about which specific one obtained. What would be the value of this market study? Intuitively again, if competitor 2 knows that competitor 1 knows the size of the market, he should be far less willing to take a bold risk and may very well fall back on a follower attitude. But this is not so simple since if competitor 1 could expect a follower attitude, he could exploit competitor 2's belief by reversing his choices (set a high price in a good market, and get 7 units, and a low price in a bad market, and get 5 units). Now, if competitor 2 could expect that competitor 1 expects a follower attitude, he could exploit competitor 1's belief . . . . Clearly, the inconsistency in this succession of mutual expectations may only be resolved using randomization. This is confirmed by the game theoretical analysis which we shall now present and interpret.

4.2. The value of information to competitor 1 and how to get it.

Assume that $p = 1/2$, then from a theoretical standpoint we know that,

$$\mathrm{EVPI}(1/2) = \mathrm{Cav}\ \bar{u}(p)\big|_{p=1/2} - \bar{u}(1/2)$$

$$= 5/11 \cdot \bar{u}(1/5) + 6/11 \cdot \bar{u}(3/4) - \bar{u}(1/2).$$

In order to understand this we shall introduce an intermediary step. Assume that perfect information is not available to competitor 1 but that the following experiment is available, $I^0 = \{i, j\}$ such that $p_i = 1/5$ and $p_j = 3/4$ with respective marginal probabilities $\gamma_i = 5/11$ and $\gamma_j = 6/11$. Moreover assume that the outcome of the experiment will be made public to both competitors. Note that at the points p_i and p_j, competitor 2 is completely indifferent between setting a high or a low price. Then, the value of this public experiment to competitor 1 is $5/11\bar{u}(1/5) + 6/11\bar{u}(3/4) - \bar{u}(1/2)$.

If competitor 1 could privately buy the experiment I^0 he always has the option to make the outcome public so that

$$\mathrm{EVI}(1/2 \mid I^0) \geqslant (5/11)\bar{u}(1/5) + (6/11)\bar{u}(3/4) - \bar{u}(1/2).$$

Can he do better?

If competitor 1 does not make the outcome public, competitor 2 is no longer indifferent between which price to set but should set the opposite price of competitor 1. Such an attitude cannot be exploited since, if competitor 1 decided to switch his prices (set a high price in case of $p_i = 1/5$, and a low price in case of $p_j = 3/4$), he would himself be worse off (for instance he would obtain $(4/5) \cdot 4 + (1/5) \cdot 3 = 19/5$, instead of $(4/5) \cdot 6 + (1/5) \cdot 1 = 25/5$, by changing from a low to a high price in case of $p_i = 1/5$). Consequently, whether or not competitor 1 makes the outcome public is irrelevant (it only makes competitor 2's problem somewhat simpler) and so

$$\mathrm{EVI}(1/2 \mid I^0) = (5/11)\bar{u}(1/5) + (6/11)\bar{u}(3/4) - \bar{u}(1/2).$$

Surprisingly enough, according to our theory, $\mathrm{EVPI}(1/2) = \mathrm{EVI}(1/2 \mid I^0)$; that is, the private value of perfect information to competitor 1 is equal to the public value of imperfect information to both competitors. This is explained as follows. What would be the public value of perfect information? Clearly this would be $(1/2)\bar{u}(0) + (1/2) \cdot \bar{u}(1) - \bar{u}(1/2)$ which is seen to be smaller than $\mathrm{EVI}(1/2 \mid I^0)$. So if competitor 1 gets perfect information, then he is no longer indifferent between making the outcome public or not. Intuitively he knows too much to make it public! Theoretically he should delete his surplus of information by putting himself back into partial ignorance. If he learns that the market is bad, he should claim that it is only bad with probability 3/4, and if he learns that it is good, he should claim that it is good only with probability 4/5. If competitor 1 cannot make his claims believed, then the only opportunity which remains is to randomize his choices according to the following table:

price	good market	bad market
high	3/11	9/11
low	8/11	2/11

Then competitor 2 will use the price set by competitor 1 as an imperfect experiment on the state of the market. Using Bayes' rule he may, for instance, derive that

$$\text{Prob(good market} \mid \text{1's price is low)} = \frac{(8/11)\cdot(1/2)}{(8/11)\cdot(1/2)+(2/11)} = 4/5$$

which, of course, is precisely what competitor 1 claimed when he set a low price. Since competitor 1 may theoretically get rid of his surplus of information using a randomized choice, it is clear (and it is also intuitive) that $\text{EVPI}(p_0) \geqslant \text{EVI}(p_0 \mid I^0)$. It remains to be seen that he cannot do better. Again competitor 2 has a strategy, involving randomization, which cannot be exploited. It is given by the following table:

price set by Competitor 1	Competitor 2 low	Competitor 2 high
low	4/11	7/11
high	5/11	6/11

The effect of this strategy is to make competitor 1 indifferent between which price to set whatever the market is (for instance if the market is good, competitor 1's expectations are $5\cdot(4/11)+6\cdot(7/11)=(62/11)$ in case of low price, and $4\cdot(5/11)+7\cdot(6/11)=(62/11)$ in case of a high price). Consequently it is not only a Bayesian best reply for competitor 2, since it optimizes 2's expected payoff conditional on competitor 1's price, but it is a reinforcement for competitor 1's own randomization. In terms of expected payoff we finally obtain $\text{EVPI}(p_0) = \text{EVI}(p_0 \mid I^0)$. If we note that the experiment I^0 is indeed the experiment whose public value is the highest for competitor 1, this gives an interesting interpretation to Proposition 3.1.

5. Further Remarks on the Value of Information

It is sometimes suggested that the value of information in a competitive situation benefits from two components: (i) the very use of information such as in a classical decision problem under uncertainty (there are more strategies available with than without information), (ii) the simple fact that the opponent believes that one has perfect information (some of these new strategies may be used only as threats).

In this section we shall first present a specific interpretation of this general statement in the context of our case study. Then, we shall show by means of another example that this second component of the value of information may turn out to be disadvantageous for the competitor who receives information, so that altogether the value of information may be negative.

5.1 The case revisited.

Assume now that competitor 1 will have perfect information with probability q ($0 \leqslant q \leqslant 1$), and that with probability $1-q$ he will remain uncertain, just as competitor 2. As usual, assume that competitor 2 knows that much. What is the value of such

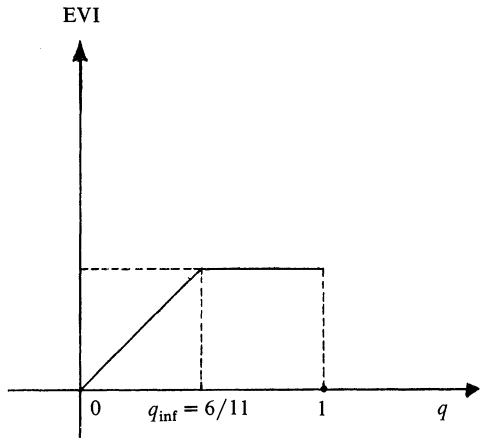

FIGURE 3. EVI as a function of the probability of having perfect information.

an experiment to competitor 1? We already know the answer for $q = 1$ (see preceding section), and $q = 0$ since then its value is of course zero. It is a simple exercise to repeat the analysis for this experiment, as done in §4, but now with three potential states of information for competitor 1; that is, the former two, bad or good market, and a new one in which he remains uncertain, with respective marginal probabilities pq, $(1-p)q$, $(1-p)(1-q)$.

For $p = 1/2$, the analysis is summarized in Figure 3: the EVI associated with this new experiment as a function of q $(0 \leqslant q \leqslant 1)$. Note that as long as q remains greater or equal to $6/11$, the EVI remains constant; as long as competitor 2 believes competitor 1 has perfect information with a high enough probability $q \geqslant q$ inf (in this example q inf $= 6/11$), this is enough to ensure the maximal benefit from the information. As a matter of fact, in many strictly competitive situations this probability will turn out to be surprisingly small, thus giving much credit to the second point of the statement mentioned at the beginning of this section.

5.2. A nonstrictly competitive situation.

Information usage may be much more complicated in nonstrictly than in strictly competitive situations. Our goal here is simply to show by means of an example that some of the new strategies available with information may represent sufficient threat to deter the opponent, to make him change his own behavior and all together to turn out to generate a negative value to the competitor having information (for more examples see [4]).

Indeed suppose that competitors 1 and 2 may engage in a mutually advantageous lottery in the following way: (i) a fair coin is tossed and not shown to the competitors, (ii) competitor 1 announces heads or tails, (iii) competitor 2, knowing competitor 1's announcement, may engage the two of them in the lottery by saying the reverse, in which case the competitor who is correct would receive four units and the other one zero, or refuse the lottery by saying the same thing as competitor 1, in which case both competitors would receive one unit each, irrespectively of which side of the coin prevails.

Since the expectation of the lottery is, $4 \cdot (1/2) + 0 \cdot (1/2) = 2$, the second competitor should engage both of them in the bet; that is, say heads if the first one says tails and vice versa. Competitor 1 may indifferently say heads or tails so that both of them have an expectation of two units. In this context, what is the value to competitor 1 if it is privately revealed to him on which side the coin fell? Clearly, if competitor 2 knows that competitor 1 knows, and if side payments are not allowed, the lottery is no longer fair. It is a dominating strategy for competitor 1 to announce the true state of the coin since if the lottery is played he receives four units for sure and if it is not played, it does not matter what he says. Consequently, competitor 2 should of course refuse to play the lottery so that finally both of them have now an expectation of one unit! Thus, the value of this private information to competitor 1 is minus one unit.

6. Discussion and Summary

In this paper we have been interested in investigating the value of information in a competitive environment. It was assumed that *if the decision maker could acquire some information, then his competitor would know that experimentation took place though he would ignore the specific outcome of the experiment.* Moreover it was assumed that the competitor would especially be aware of the acquisition of information because he could observe the decision maker's eventual change of behavior. Admittedly the analysis of such real situations would be quite complicated. The objective of the paper has merely been to present a game model of such situations in the hope that its analysis could offer some practical insights. Our main findings for strictly competitive situations may be summarized as follows:

(i) The decision maker who makes the experimentation should plan that his competitor will learn, but, since he is the one who gets the information, he can control his learning to his own advantage.

(ii) In our model this controlled learning results in the fact that the value of perfect private information is equal to the value of the public experiment which would be the most profitable for the decision maker.

(iii) This public experiment will ordinarily be an imperfect experiment because in a competitive environment uncertainty need not be disadvantageous (i.e. the payoff function may not be convex in terms of the uncertainties).

(iv) The control of the competitor's learning derived so as to keep the benefit of the uncertainties may be difficult since it could theoretically involve bluffing, the practice of which is somewhat risky.

As for nonstrictly competitive situations, no general statement could be made. Indeed an example is provided in which the value of information is in fact negative.

References

1. HARSANYI, J. C., "Games with Incomplete Information Played by 'Bayesian' Players," Parts I–III, *Management Science*, Vol. 14, Nos. 3, 5, 7.
2. HOWARD, R. A., "Information Value Theory," *IEEE Trans. on Systems Science and Cybernetics*, Vol. SSC-4 (June 1967), pp. 54–60.
3. LUCE R. D. AND RAIFFA, H., *Games and Decisions*, Wiley, 1957.
4. PONSSARD, J. P., "Information Usage in Noncooperative Game Theory," Ph.D. dissertation, EES Department, Stanford University, California (September 1971).
5. ———, "A Note on Information Value Theory for Experiments Defined in extensive Form," *Management Science* (Dec. 1975).
6. ——— AND ZAMIR, S., "Zero-Sum Sequential Games with Incomplete Information," *The Int. J. of Game Theory*, Vol. 2, No. 2, pp. 99–107 (1973).
7. RAIFFA, H. *Decision Analysis, Introductory Lectures*, Addison-Wesley, 1967.