RM-74-11

THE VALUE OF INFORMATION IN STRICTLY COMPETITIVE SITUATIONS

Jean-Pierre Ponssard

June 1974

The Value of Information in Strictly Competitive Situations*

Jean-Pierre Ponssard**

Abstract

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 non-competitive 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 strictly competitive situations.

§1 Introduction

The concept of the value of information is one of the cornerstones of decision analysis [3, 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

---

*This paper was presented at the "Seminar in Decision Theory," Beaulieu, France, December 1973, which was sponsored by the Centre d'Enseignement Supérieur du Management Public, 94112, Arcueil, France. It will be published in the proceedings of that conference.

**On leave from the Centre d'Enseignement Superieur du Management Public, 94112, Arcueil, and Groupe de Recherche en Gestion des Organisations Ecole Polytechnique, 75005, Paris, France; research scholar at the International Institute for Applied Systems Analysis, Laxenburg, Austria.

The author would like to express his thanks to Robert Winkler for extensive discussions on the subject.

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 analyses 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, 5].

In section 2 we shall define the concept of the value of information as used in decision analysis. In section 3 we shall show how the concept may be extended to strictly competitive situations. This will be illustrated by means of an example in section 4.

§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 $p_o = \{p_o^e\}_{e\varepsilon E} (p_o^e > 0; \sum_{e\varepsilon E} p_o^e = 1)$, 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 $p\varepsilon P = \{p^e | p^e \geq 0, \sum_{e\varepsilon E} p^e = 1\}$, the optimal expected payoff $\bar{u}(p)$ is then obtained as

$$\bar{u}(p) = \underset{a\varepsilon A}{\text{Max}} \sum_{e\varepsilon E} u(a,e)p^e \quad . \tag{2.1}$$

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

$$\sum_{i\varepsilon I} \gamma_i p_i = p_o \quad . \tag{2.2}$$

An experiment may equivalently be defined by a matrix $Q = \{q_{ei}\}_{e\varepsilon E, i\varepsilon I}$ in which $q_{ei} = \text{Prob}\ \{i|e\}$. One may go from one definition to the other one by means of Bayes theorem. We shall mostly use the first definition (for a practical justification of this definition see, for instance, example 1.4.3 in [7]; for further theoretical ramifications see [6]).

The expected value of the information to be revealed by the experiment I^o, $EVI(p_o|I^o)$, is then defined as the incremental gain obtained by making one's decision depend on the outcome of the experiment. Namely

$$EVI(p_o|I^o) = \sum_{i \varepsilon I} \gamma_i \bar{u}(p_i) - \bar{u}(p_o) \quad . \qquad (2.3)$$

As a special case the expected value of perfect information, $EVPI(p_o)$, obtained by the experiment which would reveal the state of nature, is such that (writing p_e for the probability vector such that $p_e^k = 0$ for all $k \neq e$ and $p_e^e = 1$)

$$EVPI(p_o) = \sum_{e \varepsilon E} p_o^e \bar{u}(p_e) - \bar{u}(p_o) \quad . \qquad (2.4)$$

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

In the remaining of this section we shall prove some simple properties suggested by (2.3), (see [6] for a full discussion of these properties). This will also allow us to introduce the technical apparatus needed subsequently.

Denote by P_I the smallest convex subset of P which contains the vectors $\{p_i\}_{i \varepsilon I}$ and by $\underset{P_I}{Cav}\, f(p)$ the minimal concave function* greater or equal to f(p) on P_I, in which f(p) is any real-valued

*g(p) is a concave function on P if and only if for all p_1 and p_2 and all $\lambda \varepsilon (0,1)$: $g(\lambda p_1 + (1-\lambda) p_2) \geq \lambda g(p_1) + (1-\lambda) g(p_2)$.

continuous function on P. Let $\operatorname{Cav}_{P_I} f(p)\Big|_{p_o}$ stand for the value of the function $\operatorname{Cav}_{P_I} f(p)$ at p_o.

Proposition 2.1. For any experiment I^o and any $p_o \varepsilon P_I$, the value of information, $EVI(p_o\ I^o)$, satisfies

$$EVI(p_o|I^o) \leq \operatorname{Cav}_{P_I} \bar{u}(p)\Big|_{p_o} - \bar{u}(p_o) \quad . \tag{2.5}$$

Proof. The inequality follows directly from the definition of Cav_{P_I} and from (2.2) and (2.3). ||

Corollary 2.2. If the set of vectors $\{p_i\}_{i\varepsilon I}$ are linearily independent then (2.5) is an equality.

Proof. Let $\Lambda^o = \{\lambda = (\lambda_i)_{i\varepsilon I} | \lambda_i \geq 0 \ , \ \sum_{i\varepsilon I} \lambda_i = 1 \ , \ \sum_{i\varepsilon I} \lambda_i p_i = p_o\}$.

Since $\bar{u}(\cdot)$ is a convex function on P_I, its concavification may be expressed as

$$\operatorname{Cav}_{P_I} \bar{u}(p)\Big|_{p_o} = \operatorname{Max}_{\lambda\varepsilon\Lambda^o} \sum_{i\varepsilon I} \lambda_i \bar{u}(p_i) \quad .$$

But the set of vectors $(p_i)_{i\varepsilon I}$ are linearily independent, hence Λ^o contains only one point, namely $\gamma = \{\gamma_i\}_{i\varepsilon I}$, the probability distribution on the set $\{p_i\}_{i\varepsilon I}$ associated with experiment I^o. Thus

$$\operatorname{Cav}_{P_I} \bar{u}(p)\Big|_{p_o} = \sum_{i\varepsilon I} \gamma_i \bar{u}\ (p_i) \quad . \qquad ||$$

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 $P = \{p = (p^1, p^2) | p^1 \geq 0, p^2 \geq 0, p^1 + p^2 = 1\}$ ($\bar{u}(\cdot)$, a convex function, is piece wise linear since the set of actions A is finite). Let $p_o \varepsilon P$ be the a priori probability distribution on E and the experiment $I^o = \{i,j\}$ be defined by two possible a posteriori probability distribution on E, $p_i \varepsilon P$ and $p_j \varepsilon P$, with marginal probabilities γ_i and γ_j respectively (recall that for consistency we have $\gamma_i p_i + \gamma_j p_j = p_o$). Then the information value analysis is completely described by the following graph.

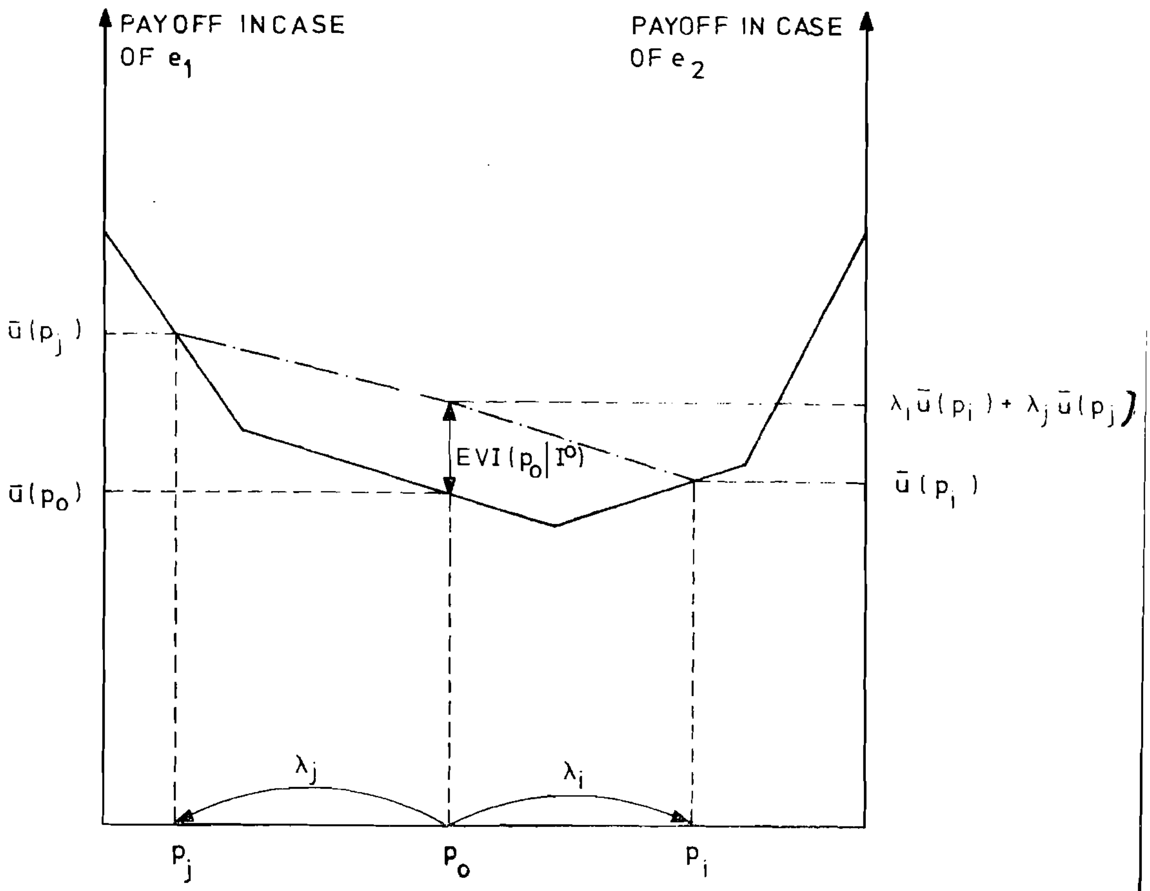

Finally note a simple result as a direct specification of Corollary 2.2.

Corollary 2.3. The value of perfect information may by expressed as,

$$\text{EVPI}(p_o) = \underset{P}{\text{Cav}}\ \bar{u}(p)\Big|_{p_o} - \bar{u}(p_o) \quad . \tag{2.6}$$

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 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

$$(a \in A) \quad (b \in B) \qquad u(a,b,e) + v(a,b,e) = c(e) \quad .$$

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

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

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

$$\bar{u}(p_o) = \max_{a \in A} \min_{b \in B} \sum_{e \in E} p_o^e u(a,b,e) \quad , \tag{3.1}$$

and 2's optimal expected payoff is

$$\bar{v}(p_o) = \underset{a \epsilon A}{\text{Min}} \ \underset{b \epsilon B}{\text{Max}} \ \sum_{e \epsilon E} p_o^e v(a,b,c) \ . \qquad (3.2)$$

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, <u>assuming that the other one will know</u> 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 [5, theorem 1,page 101]. In the context of this paper the result appears as an extension of corollary 2.2.

<u>Proposition 3.1</u>. In a strictly competitive sequential situation the value of perfect information to competitor 1 may be expressed as

$$\text{EVPI}(p_o) = \underset{P}{\text{Cav}} \ \bar{u}(p) \Big|_{p_o} - \bar{u}(p_o) \ . \qquad (3.3)$$

Insights provided by this result and their interpretations 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 significent 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, $a \varepsilon A$ for 1 and $b \varepsilon B$ for 2, for a new product. Moreover, suppose that the size of the market, $e \varepsilon E$, 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 look as follows:

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)

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)

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

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

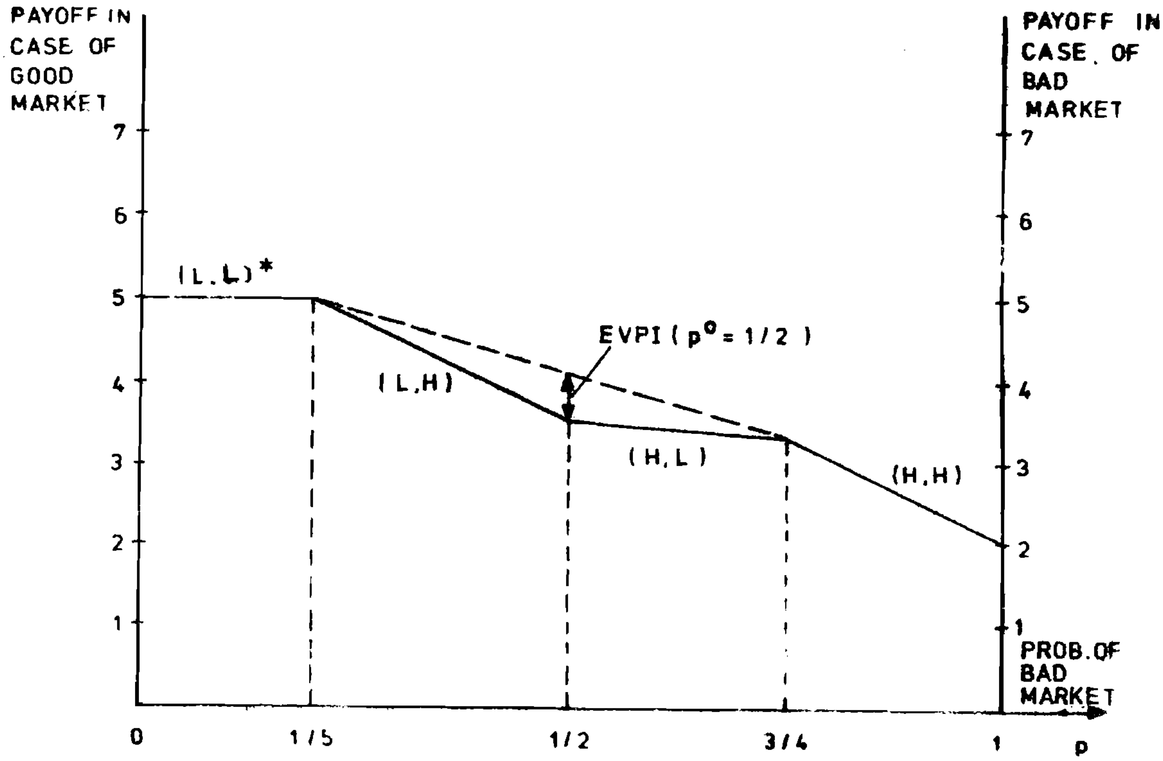

Figure 2. 1's optimal expected payoff $\Sigma\ \bar{u}(p)$

---

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

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 1 unit if 1's expectation turns out to be right but he wins 4 units if 1's expectation turns out to be wrong). We shall concentrate our analysis in the case of high uncertainites ($\frac{1}{5} \leq p \leq \frac{3}{4}$).

Suppose now that competitor 1 may order a market study and thus obtain perfect information (competitor 2 would know that competitor 1 ordered a market study though he would not know the result). 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 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,

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

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

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: $I^o = \{i, j\}$ 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 precisely indifferent between setting a high or a low price. Anyhow the value of this public experiment to competitor 1 is

$$\frac{5}{11}\, \bar{u}(1/5) + \frac{6}{11}\, \bar{u}(3/4) - \bar{u}(1/2) \quad .$$

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

$$\text{EVI}(1/2|I^o) \geq \frac{5}{11}\, \bar{u}(1/5) + \frac{6}{11}\, \bar{u}(3/4) - \bar{u}(1/2) \quad .$$

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

$$\frac{4}{5} \cdot 4 + \frac{1}{5} \cdot 3$$

instead of

$$\frac{4}{5} \cdot 6 + \frac{1}{5} \cdot 1$$

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

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

Suprisingly enough, according to our theory,

$$EVPI(1/2) = EVI(1/2|I^{o}) \ ;$$

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 $\frac{1}{2}\,\bar{u}(0) + \frac{1}{2}\,\bar{u}(1) - \bar{u}(1/2)$, which is seen to be smaller than $EVI(1/2|I^{o})$. 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

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

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

$$EVPI(p_o) \geq EVI(p_o | I^o) .$$

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 2	
Competitor 1	low	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

$$5 \cdot \frac{4}{11} + 6 \cdot \frac{7}{11} = \frac{62}{11}$$

in case of low price, and

$$4 \cdot \frac{5}{11} + 7 \cdot \frac{6}{11} = \frac{62}{11}$$

in case of a high price). Consequently it is not only a Bayesian strategy for competitor 2, since it optimizes his 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

$$EVPI(p_o) = EVI(p_o | I^o) \quad .$$

If we note that the experiment I^o is indeed the experiment whose public value is the highest for competitor 1, this gives an interesting interpretation to Proposition 3.1.

§5 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 acquision 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 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 their 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 the competitive environment uncertainty need not be disadvantageous (i.e. the payoff function may not be convex in terms of the uncertainties); and

(iv) the control of the competitor's learning derived so as to keep the benefit of the uncertainties may be a difficult practical matter.

References

[1] Harsanyi, J.C., "Games with Incomplete Information Plays by 'Bayesian' Players," Parts I, II, 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 H. Raiffa, Games and Decision. New York, Wiley, 1957.

[4] Ponssard, J.P., "Information Usage in Non-Cooperative Game Theory," Ph.D. dissertation, EES Department, Stanford University, California, September 1971.

[5] Ponssard, J.P. and S. Zamir, "Zero-Sum Sequential Games with Incomplete Information," The Int. J. of Game Theory, 2, No. 2 (1973), 99-107.

[6] Ponssard, J.P., "A Note on Information Value Theory for Experiments Defined in Extensive Form," Working Paper No. WP-74-21, June 1974, IIASA, Laxenburg, Austria.

[7] Raiffa, H. Decision Analysis, Introductory Lectures. Reading, Mass., Addison-Wesley, 1967.