Endogenous Private Information Structures*

Sjaak Hurkens
Department of Economics and Business
Universitat Pompeu Fabra
Ramon Trias Fargas 25-27
08005 Barcelona, Spain

Nir Vulkan
Economics Department
Bristol University
8 Woodland Road
Bristol BS8 1TN, UK

December 1999

Abstract

Many models in the economics literature deal with strategic situations with privately informed agents. In those models the information structure is assumed to be exogenous and common knowledge. In many applications information gathering is one of the strategic options available to agents. We formally incorporate this option into the game and the information structure will arise endogenously. We ask whether models with exogenous information structures, and the results they provide, are robust with respect to this endogenization. We show that any Nash equilibrium of the game with information acquisition induces a Nash equilibrium in the corresponding game with an exogenous structure. The same is not always true when 'Nash equilibrium' is replaced by 'sequential equilibrium' but we provide sufficient conditions on the structure of the game for which this is true. Moreover, we characterize the (sequential) Nash equilibria of games with an exogenous information structure that can arise as a (sequential) Nash equilibrium of a game with endogenous information acquisition,

Journal of Economic Literature Classification Numbers: C72, D82, D83.

Keywords: Information acquisition, private information, equilibrium.

*This paper is a substantial revision of 'Endogenous Information Structures' that was circulated in April 1999. We thank seminar audiences at UPF and UAB for their comments. Hurkens gratefully acknowledges financial support from CIRIT, Generalitat de Catalunya, grant 1997SGR 00138 and DGES grant PB96-0302. Vulkan gratefully acknowledges financial support from the EPSRC, grant no. M07052 and from the ESRC grant no. R000222642.

1 Introduction

Modern economic theory emphasizes the importance of private information in industrial organization, principal-agent relationships, bargaining, contract theory, auctions and financial markets. In those models the private information of players is an *exogenous* part of the model. Such models explain what will happen *given* a certain private information structure, but do not explain where that information structure *comes from.*[1] For some economic situations assuming an exogenous private information structure may be more sensible than for others. For example, in Akerlof's famous Lemons market it seems to make good sense to assume that owners of a used car have private information about the quality of their car, since they are supposed to have used this car in past years. On the other hand, in a principal-agent model of procurement it is less compelling to assume that the agent has private information about his cost as this cost would depend on the requirements of the project imposed by the principal. Hence, the agent will have to process those requirements and make estimates of input costs, *i.e.* the agent will have to acquire information about his cost. The private information structure of the agent would thus be endogenous. Similar stories apply for bidders in an auction, firms and entrants in imperfect competition models, and speculators in stock markets.

Now, continuing with the procurement example, one might argue that the information about cost can be much easier acquired by the agent than by the principal, and that it is obviously in the interest of the agent to do so at the (relative) small effort it takes him. This argument would then seem to justify the use of the model with an exogenously given privately informed agent. Similar arguments would of course then apply in similar situations.

However, some recent papers have shown that endogenizing the information structure leads to surprising but not necessarily unintuitive results and, more importantly, to new insights in varying fields in economic theory. In the case of principal-agent theory Kessler (1998) has shown that the agent may prefer to stay uninformed, thereby crucially changing the incentive compatibility and individual rationality constraints of the principal's problem. Cremer and Khalill (1992) show that the principal may offer contracts that induce the agent to stay uninformed and that those contracts will be very different from

[1] Notable exceptions are Vives (1988), Li *et al.* (1987), Hwang (1993, 1995), Hauk and Hurkens (1997), and Ponssard (1979) that consider information acquisition in oligopolies, Lee (1982), Milgrom (1981), Matthews (1984), and Persico (1997) who consider the incentives for information acquisition in auctions and Cremer and Khalil (1992, 1994), Cremer *et al.* (1998), and Kessler (1998) that consider principal agent relationships where the agent is initially uninformed but can acquire information, before or after the principal has offered him a contract. Grossman and Stiglitz (1980) initiated research on the incentives to acquire costly information on the value of a stock in financial markets. Hurkens and Vulkan (1996, 1997) considered information acquisition by potential entrants.

those that a principal would offer in case the uninformed agent would not have the possibility to gather information. In the context of auctions Persico (2000) shows that the incentives for bidders to gather information are different in first- and second-price auctions so that the revenue equivalence theorem breaks down. Finally, Hurkens and Vulkan (1997) show that the number of entrants in a new market with uncertain demand when firms' private information about demand is exogenously given need not coincide with the number of entrants when firms have to decide whether or not to become informed. Surprisingly, this holds even though the private information structures that arise in the endogenous model are identical to the ones assumed in the exogenous one.

It is not unlikely that this list of examples can be extended with applications from other fields. That is however not the main purpose of this paper. The aim of this paper is to investigate, at an abstract and general level, which type of games are robust and immune with respect to the endogenization of the private information structure and which type of games are susceptible to this endogenization. Games from the first class can thus be analysed using an exogenous private information structure even when they in fact describe economic situations where information gathering seems a natural part of players' strategic options. Games that belong to the second class, however, will have to be considered with care as they are possibly affected by having endogenous rather than exogenous information structures. Applications in the line of the three mentioned examples will have to be found in this class.

Our investigation of robustness starts out of two important questions. First, can any exogenous private information structure arise endogenously? Second, does behaviour in the game depend on whether the private information structure is exogenously given or is endogenously determined? In fact, if the answers to these questions were 'yes' and 'no' respectively, we would be fully justified to use models with exogenous private information structures. From our examples discussed before we know that in fact the answers will not always be 'yes' and 'no', and we need to investigate in which cases the answers may differ.

At first sight the answer to the first question would seem to be 'yes'. Namely, it suffices to choose the costs for different information structures such that the one information structure we want to 'explain' has the highest value, *net of costs*. However, one is not completely free in the choice of costs: any reasonable information cost function should be (weakly) increasing, *i.e.* more information should cost more. A more fundamental problem is that the notion of 'value of information' in a game is rather vague. (See also Neyman (1991).) How much a player can improve his payoff by obtaining some piece of information will depend on how the other players in the game will behave and respond. Hence, the value of information in a game (if we can even speak about such a concept) must be closely linked to (equilibrium) behaviour in the game. Clearly, existence of multiple equilibria and issues of equilibrium refinements will affect and possibly complicate our task. As a matter of fact, this issue will turn out to be very important.

The answer to the second question would seem to be 'no', especially since we will be interested in equilibrium behaviour. In equilibrium of the game with endogenous information acquisition each player 'knows' how much information other players acquire, similar to the case of games with exogenous private information structures. There is a small but important difference though. Namely, a player will have beliefs about how much information others acquired, and in equilibrium these beliefs must be correct. However, *off* the equilibrium path a player may hold incorrect beliefs which may affect his behaviour *off* the equilibrium path. Of course, behaviour *off* the equilibrium path will in fact determine the behaviour *on* the equilibrium path. It is therefore not obvious at all that the answer to this second question is always 'no'. In fact, the entry example discussed before constitutes a counterexample. Since beliefs seem to play a role here the issue of equilibrium refinements (such as sequential equilibrium) again turns up.

We conclude that we cannot answer the above two questions without fully analyzing and comparing equilibrium behaviour in games with exogenous and endogenous information structures. That is what we will do in the paper. We will consider equilibrium behaviour in the game with endogenous information acquisition and compare that to equilibrium behaviour in the game with an exogenous information structure, where the information structure is in fact the one that endogenously arises in the first game. In the other direction, we will consider equilibrium behaviour in a game with an exogenously given private information structure and we ask ourselves whether this behaviour and this information structure can prevail in some equilibrium of the endogenous information acquisition game. We also investigate the importance of equilibrium refinements.

Let us now outline the results obtained in this paper. We will focus on the class of games where information about fundamentals (*i.e.* moves of Nature that affect the payoffs but not the strategic options of the players) is acquired before players enter into a strategic situation. Hence, there is an information gathering stage followed by a game playing phase.[2] Information costs are exogenously specified but may differ amongst player to reflect the asymmetry in the underlying game. Since the actions of a player are chosen after information acquisition decisions are taken, they cannot influence those decisions. By means of an example we show that the information acquisition game may have only equilibria in which players randomize over several possible private information

[2]This excludes the case analyzed by Cremer and Khalil (1992) where the principal offers a contract before the agent may gather information. It also excludes the case where a player can get information about *unobserved actions* chosen by other players. See Perea y Monsuwé (1997) for a setting in which players buy information about actions chosen by other players.

structures. For obvious reasons of comparability we are thus forced to restrict attention to equilibria of the endogenous information acquisition games where players use a pure strategy in the information gathering phase. It should be clear that to any of such equilibria corresponds a game with an exogenous information structure, together with some strategy profile in that exogenous game.

Our first result says that any 'pure information' equilibrium of the endogenous game induces an equilibrium in the corresponding exogenous game. The result is simple and intuitive but also important. Namely, it follows that games with exogenous information structures that have a unique Nash equilibrium are quite robust to endogenization of the information in the sense that behaviour in the game does not depend on whether the information structure is exogenously given or endogenously determined. (Of course, as seen in the principal-agent examples, it might be true that the given information structure could not have arised endogenously for any reasonable specification of information costs.) We then consider equilibrium refinements, and in particular we consider sequential equilibria. We show that any sequential pure information equilibrium of the endogenous information acquisition game induces a sequential equilibrium in the corresponding game with exogenous structure if (i) the information acquisition decisions are perfectly observed, or (ii) the game playing phase only involves simultaneous moves. A counterexample illustrates that if neither of these conditions is satisfied a sequential equilibrium of the endogenous information acquisition game need not induce a sequential equilibrium in the exogenous game. In fact, the same example will do for many other refinements as well. A consequence of this is that games with exogenous information structures that have a unique (plausible) sequential equilibrium but multiple Nash equilibria may be affected by the endogenization of information.

We then turn our attention to equilibrium behaviour in games with exogenous information structures and ask ourselves whether this behaviour and information structure can also arise in the endogenous information acquisition game for some reasonable assumptions on information costs. If we do not exclude that information is available for free then any Nash equilibrium of a game with an exogenous private information structure may arise as a Nash equilibrium of the game with endogenous information acquisition. When information acquisition is secret and information is not free of cost then this only holds in case all information is used in the equilibrium of the exogenous game under consideration. Hence, pooling in signalling games or bunch-bidding in auctions are excluded in this case. On the other hand, when information acquisition is observed and we restrict attention to sequential equilibria then the result holds only if information does not 'hurt', that is, if no player is better off (in any equilibrium) with less information.

The rest of the paper is set up as follows. Section 2 models private information and defines two modes of information acquisition, *secret* and *private*. We start the analysis in section 3 by means of an example to illustrate the difficulties that may arise in comparing endogenous and exogenous information structures. We then consider equilibria of the endogenous games and analyze the behaviour it induces in the corresponding game with an exogenous information structure. The final part of this section goes in the other direction, that is, we consider an equilibrium of some exogenous game and investigate when the same information structure and behaviour may arise when information acquisition is endogenous. Section 4 summarizes the results and hints at some possible applications for future research.

2 The Model

There are n players involved in a strategic interaction represented by a game form G. This game form describes the actions that players may take, and the order in which they take them. It includes information sets, indicating that players may not know which actions have been chosen before. It is not, however, a fully defined extensive form game because there are no payoffs associated to the end points of the tree. The payoff functions to evaluate outcomes depend on the realized value of the state of Nature. We denote by Ω the finite set of states of Nature and will write ω to denote a generic state. We denote by S_i the finite set of pure strategies of player i in G (which does not depend on the state of Nature) and write $S = \prod_i S_i$ as the product set of strategy profiles. We denote by $u_i^\omega(s)$ the payoff obtained by player i when strategy profile s is chosen and the state of Nature is ω.

The players initially hold a common prior about the likelihood of each state of Nature which is represented by a probability distribution ρ on Ω. However, players can either endogenously improve their information or will be exogenously endowed with better information. The information structure of a player is given by some partition P of Ω. Hence, $P = \{P_1, ..., P_k\}$ where $\cup_{i=1}^k P_i = \Omega$ and $P_i \cap P_j = \emptyset$ for all $i \neq j$. We write $P(\omega)$ for the element of P that contains ω. The interpretation of the partition is that with information structure P a player can distinguish between two states ω and ω' if and only if $P(\omega) \neq P(\omega')$. We say that partition P is finer (more informative) than partition P' if $P(\omega) \subset P'(\omega)$ for all $\omega \in \Omega$ (with strict inclusion for some ω).

Each player i will either be endogenously or exogenously endowed with some information structure P^i. Let $P = (P^1, ..., P^n)$. We will denote the game with the exogenous information structure P by G^P. This is a standard game with private information. We will compare the equilibria of such games with those of the super-game where information structures are endogenously determined. This super-game has two stages. In the first stage all players choose simultaneously one of the feasible information structures. Not all information structures need to be feasible but we assume that $P^{\text{no}} = \{\Omega\}$ (no information acquisition) and $P^{\text{full}} = \{\{\omega\} : \omega \in \Omega\}$ (full information) are. Information acquisition is costly and costs are player specific. Player i needs to pay $c(i, P)$ to acquire partition P, where $c(i, P) \geq c(i, P')$ if P is more informative than P'. When information is not freely available the inequalities will be strict. For convenience, we assume $c(i, P^{\text{no}}) = 0$.

Information acquisition can either be *secret* or *private*. In the super-game where information is acquired secretly, denoted by Γ^s, the choice of partition P^i by player i is not observed by players $j \neq i$. A pure super-game strategy for player i in this game is a pair (P^i, σ_i), where P^i is the partition chosen and $\sigma_i : \Omega \to S_i$ maps states of Nature into strategies of G. Of course, the mapping σ_i must be measurable with respect to P^i, i.e. $\sigma(\omega) = \sigma(\omega')$ if $P^i(\omega) = P^i(\omega')$. On the other hand, in the super-game where information acquisition is private, denoted by Γ^p, the choice of partition P^i by player i is observed by all players. Hence, a pure super-game strategy for player i is now a pair (P^i, τ_i), where P^i is again the partition chosen by i and where $\tau_i(\omega, P) \in S_i$ denotes the strategy chosen in G when the state of Nature is ω and the (total) information structure is P. Again, the mapping τ_i must be measurable with respect to P^i.

We will be interested in comparing the results of the endogenous information acquisition games Γ^s and Γ^p with those obtained in games with an exogenous information structure, G^P. Since the strategy spaces in those three types of games are different, we need to make precise how we will compare the equilibrium results. First, in order for the comparison to make any sense we must restrict attention to those equilibria of the endogenous information acquisition games that use pure strategies in the first stage, i.e. where players choose a unique information structure and do not randomize. As we will see in the next section, this restriction will exclude some games from consideration as they only have equilibria in which players randomize in the information acquisition stage.

Because of the sequential structure of the information acquisition games we will sometimes want to restrict attention to sequential equilibria so that our results do not rely on incredible out of equilibrium threats or beliefs. Recall that a sequential equilibrium is a pair (μ, σ) where strategy σ is optimal given the beliefs μ, and the beliefs are consistent, i.e., $(\mu, \sigma) = \lim_n(\mu_n, \sigma_n)$, where σ_n is a completely mixed strategy profile and μ_n are beliefs determined by Bayes' rule based on σ_n. Finally, we will compare the strategies induced in the second stage by the sequential equilibria of the information acquisition games with the (sequential) equilibria of the games with an exogenous information structure, where this information structure is the one that endogenously emerged in the equilibrium at hand.

In order to compare sequential equilibria of the three models, we explain how the information sets in those models correspond to each other. In an information set I_i^s for player i in Γ^s, the player knows its own information structure P^i, but does not know the information structure of the other players, P^{-i}. Moreover, it can distinguish states ω and ω' if and only if $P^i(\omega) \neq P^i(\omega')$. Hence I_i^s is simply the union over all possible information structures P^{-i} of the other players of the corresponding information sets in the exogenous games $G^{P^i,P^{-i}}$. On the other hand, in an information set of player i in Γ^p, player i knows the information structures of players $j \neq i$. Hence, each information set in Γ^p corresponds to an information set in G^P for exactly one information structure P. This allows us to identify belief assessments in Γ^p and G^P, and to imbed belief assessments in G^P into the belief assessments in Γ^s as follows: For some belief assessment μ in G^P, let $\mu^s_{|P}$ denote the assessment in Γ^s at the corresponding information set that assigns zero probability to the event that any player j chose an information structure different from P^j. In other words, player i believes that the information structure is P (with probability one) and his beliefs about which point in the information set has been reached is then given by μ.

3 Results

3.1 Comparing endogenous and exogenous information structures

Suppose that Nature determines which of the two bimatrix games of Fig. 1 is going to be played, I or II. Game I is picked with probability 1/2.

	L	R
T	3,1	6,0
M	0,0	0,1
B	4,4	5,5

I

	L	R
T	0,0	0,1
M	3,1	6,0
B	4,4	5,5

II

Fig. 1.

Only player 1 can secretly learn whether game I or II is chosen, at some cost $c \in (0,1)$. Note that learning and not using the information is a strategy strictly dominated by not learning and playing B. Also, learning and playing MT (*i.e.* M in game I, T in game II), MB, TB, BM, or BT are dominated. The only undominated strategies for player 1 are (not learn, play B) and (learn, play TM). Of course, if player 1 plays B, player 2 prefers to choose R, against which player 1 prefers TM. But if player 1 plays TM, player 2 prefers to choose L, against which it is optimal for player 1 to choose B. Hence, the game with secret endogenous information acquisition has no pure Nash equilibrium and there is a unique mixed Nash equilibrium: player 1 randomizes equally between TM and B, while player 2 chooses L with probability (1-c)/2.

Note that the information structure that endogenously arises is random since player 1 can be informed or uninformed.[3] This result is somewhat reminiscent of the principal-agent models by Kessler (1998) and Cremer and Khalil (1992) where it is shown that the agent will (with positive probability) be uninformed about his cost. Neither the model with an exogenously informed player 1, nor the model with an exogenously uninformed player 1 can be generated from endogenous and secret information acquisition, as long as information cost is positive but small.[4] Because of these complications in comparing endogenous and exogenous information structures we will henceforth restrict attention to those equilibria of the information acquisition games where players use a pure strategy in the first stage.

3.2 Endogenous Information Acquisition Equilibria

In this subsection we consider pure information equilibria of games with endogenous information acquisition. In particular, we investigate the behaviour induced by such equilibria in the game playing phase. We will consider both secret and private information acquisition games and will also discuss the role of equilibrium refinements. We start of with ordinary Nash equilibria.

Theorem 1 *(i) Let (P, σ) be a Nash equilibrium of the secret information acquisition game Γ^s. Then σ is a Nash equilibrium of G^P.*

(ii) Let (P, σ) be a Nash equilibrium of the private information acquisition game Γ^p. Then $\sigma(P)$ is a Nash equilibrium of G^P.

Proof. We just prove (i) as the proof of (ii) is along the same lines.

[3]This could be interpreted as corresponding to a very particular game in which three types of player 1 exist, type I, type II, and a third uninformed type with exogenous probabilities 1/4, 1/4, and 1/2, respectively. Note that these probabilities do not depend on c, (as long as $c \in (0,1)$) but on the payoffs of player 2. In the game with three types a unique equilibrium exists: player 1 chooses T (if type I), M (if type II) and B (if type III), while player 2 plays L with probability 1/2 (since the cost of information is sunk).

[4]For values of $c > 1$ information acquisition is dominated and we would have support for the model with an uninformed player 1, whereas if information is costless ($c = 0$) it is an equilibrium of the secret information acquisition game to acquire information and then use the information with probability a half, while player 2 uses both strategies with equal probability.

Suppose not. Then some player strictly prefers to deviate from σ in G^P. But then this same player would also like to deviate in the game playing phase of Γ^s, contradicting the assumption that (P, σ) be a Nash equilibrium of Γ^s. □

The result is straightforward but has important consequences. Namely, it follows that a game with an exogenous information structure that has a unique Nash equilibrium is at least partly robust with respect to the endogenization of the information structure in the sense that behaviour in the game playing phase does not depend on whether the information structure is exogenous or endogenous. What of course remains to be seen is whether that information structure would arise endogenously. We will come back to this issue in the next subsection.

The reason that the result is straightforward is that we consider any Nash equilibrium. That is, we only ask whether each player's action is optimal given other players' actions. Actions off the equilibrium path of play are always optimal, while actions on the equilibrium path of play are, by definition, optimal. Of course, some actions off the equilibrium path may seem unreasonable or incredible. Hence, it is natural to consider subgame perfect or, more generally, sequential equilibria.[5] One might guess that Theorem 1 would break down when replacing 'Nash' by 'sequential' as players will make inferences about other players' informedness in the game with endogenous information acquisition whereas they do not in a game with exogenous information structure. It could be that in the former game some action off the proposed path of play can be 'rationalized' by some beliefs about informedness while no such rationalization exists in a game where the information structure is given and common knowledge. As Theorem 2 below shows, it turns out that this argument can only apply in the case of games with sequential moves and when information is acquired secretly. Namely, in the case of private information acquisition the information structure will become common knowledge even though it arises endogenously. Furthermore, in the case of secret information acquisition beliefs about informedness only play a role when actions are taken after beliefs have been revised. In the case of simultaneous move games belief revision can thus not play any role!

Theorem 2 *(i) Let (P, σ) be a sequential equilibrium of the secret information acquisition game Γ^s. Then σ is a sequential equilibrium of G^P if G is a simultaneous move game.*

(ii) Let (P, σ) be a sequential equilibrium of the private information acquisition game Γ^p. Then $\sigma(P)$ is a sequential equilibrium of G^P.

Proof. (i) By Theorem 1(i) we know that σ is a Nash equilibrium of G^P. Since G^P has only simultaneous moves it is sequential as well (where beliefs are determined by

[5]The game with information acquisition may have no or few proper subgames as Nature moves first.

equilibrium strategies and moves of nature).

(ii) By definition of sequential equilibrium (as G^P is a 'subgame' of Γ^p). □

In order to see that Theorem 2(i) is, in general, not true for sequential games consider the following situation. Nature chooses between game I and game II with equal probability.

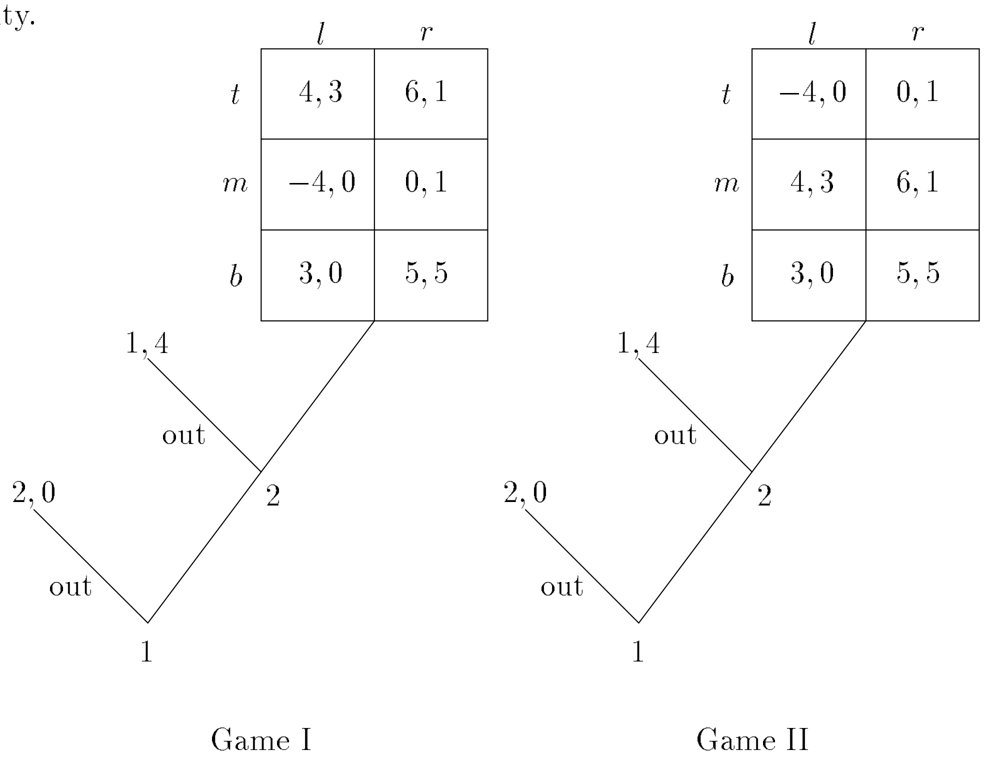

Fig. 2

We will show that the secret information acquisition game has a unique Nash equilibrium outcome where player 1 does not learn and chooses out. By Theorem 1(i) we know that this must also be a Nash equilibrium outcome in the game where it is common knowledge that player 1 is not informed. However, the only plausible (sequential) equilibrium in that game is where players enter the subgame and obtain (5,5). Hence, the sequential equilibrium of the secret information acquisition game does not induce a sequential equilibrium in the corresponding game. As we will see shortly, the reason is that in order to support the sequential equilibrium of Γ^s player 2 should believe that player 1 is informed when player 2 (unexpectedly) gets to move.

Let us now analyse the game. Player 1 can decide to learn the outcome of Nature's move at cost $c \in (0,1)$. Then he must decide to end the game (yielding gross payoff vector

(2,0)) or give the move to player 2. If given the move, player 2 can decide to end the game (yielding gross payoff vector (1,4)) or continue the game. If the game continues both players engage in a 3×2 bimatrix game. The payoffs in this game depend on whether Nature has chosen I or II. The structure of the payoffs in this 'subgame' is such that if player 1 is informed he has a strictly dominant strategy: play t if Nature choose I, choose m if Nature choose II. Player 2 will then want to play l in this subgame, yielding him a payoff of 3 (which is less than 4 which he can obtain by choosing out if given the move). If player 1 is not informed his strictly dominant strategy (in the subgame) is to play b. In that case player 2 will want to play r, yielding him a payoff of 5.

Hence, if player 2 expects that player 1 is informed, Then he will choose to end the game when he gets the chance, which in turn means that player 1 will prefer to end the game himself and obtain 2 instead of 1. Foreseeing this to happen he will thus not obtain any information. The strategy profile (NL, out, out) thus constitutes a Nash equilibrium. There are more Nash equilibria since player 2 choosing to continue the game with low probability leaves him indifferent (since he does not get to move) while it would still be optimal for player 1 to end the game immediately. Of course, when player 2 gets the move he is willing to randomize between ending and continuing the game (i.e. continuing with r as l is strictly dominated by out) only if he assigns probability 1/4 to the event that player 1 is informed . (Taking for granted that he expects an informed player to play tm and an uninformed player to choose b.) Note that for $c < 1/4$ the strategy for player 1 (NL, b) is strictly dominated by some mixture between (L, tm) and (NL, out). Hence, attaching a positive probability to the event that player 1 is uninformed when player 2 is given the move corresponds to player 2 believing that player 1 has played a strictly dominated strategy. An argument of forward induction would thus be sufficient to select our previous pure equilibrium as the unique plausible one. In this unique plausible equilibrium player 1 does not get informed but to sustain this equilibrium player 2 must believe (off the equilibrium path) that player 1 is informed.

On the other hand, if player 2 expects player 1 to be uninformed, he should expect player 1 to play b in the subgame, and he will thus choose to continue the game and play r. However, then player 1 would like to (secretly) learn Nature's move and give the move to player 2. We conclude that in the game with costly and secret information acquisition the unique equilibrium outcome is that player 1 does not learn and chooses out.

Consider now the game where it is common knowledge that player 1 is uninformed. Of course, also in this game it is a Nash equilibrium for both players to choose out. However, the only reasonable (subgame perfect) equilibrium is of course ((in, b),(in, r)).

This example is in the same spirit as our[6] application to entry in a new market. We assume there that information acquisition is secret and since firms compete in quantities after having taken entry decisions, that application has a sequential structure that makes revising beliefs possible and important. There exists a sequential equilibrium in which firms become perfectly informed about demand but less entry takes place than in a model where it is common knowledge that firms are perfectly informed. Unexpected entry is believed to be caused by uninformed entrants which makes the 'incumbent' act more agressively, thereby deterring entry alltogether.

3.3 Exogenous Information Structures

In this subsection we consider games with exogenous information structures and their (sequential) equilibria and ask ourselves whether the same information structure and behaviour could arise with endogenous information acquisition when information is costly. It will be instructive to consider some examples before coming to the formal results.

Nature determines which of the two bimatrix games of Fig. 3 is going to be played, I or II. Game I is picked with probability 2/5.

	L	R
T	4,2	1,3
B	2,2	0,0

I

	L	R
T	2,2	0,0
B	4,2	1,3

II

Fig. 3.

If neither player knows which game is played, both players have a dominant strategy and the outcome will be (B,L), yielding an expected payoff of (16/5,2). However, if player 1 observes which matrix is chosen, it is a dominant strategy to play T in game I and B in game II. If player 2 knows that player 1 observes the outcome of the move of Nature, he will then play R and the resulting payoff vector will be (1,3). Player 1 is better off in the game where it is common knowledge he does not observe Nature's move than in the game where it is common knowledge that he does. This is an example where information hurts the player possessing it. In the unique sequential equilibrium of the game with private information acquisition player 1 will thus choose (commit) not to

[6] Hurkens and Vulkan (1997).

obtain information, and the outcome will be (B,L) for any nonnegative cost of acquiring information. If information cost is small, the private information game possesses a *Nash* equilibrium in which player 1 observes Nature's move and player 2 threatens to play R in any case (that is, whether or not player 1 observes Nature's move or not). We conclude that in this example the information structure where player 1 is informed can arise in a Nash equilibrium of the private information acquisition game but not in a sequential equilibrium because information hurts.

In the case of secret information acquisition, however, player 2 cannot condition his action on whether player 1 learned or not. He will thus choose the same (mixed) strategy in both games (as he cannot distinguish them). Whatever this strategy of player 2 may be, for low enough information costs c player 1 will choose to acquire information and choose T in game I and B in game II. Player 2 will foresee this in any equilibrium and play R and the final payoffs will be $(1-c, 3)$.

We can draw some general conclusions from this example. In the case of secret information acquisition players cannot condition their strategies in the game playing phase on the information acquisition decision of any player. Hence, when making the information acquisition decision a player takes the strategies of others in the game playing phase as given and can actually calculate the value of information. This value is always non-negative and if it is strictly positive he will acquire the information for low enough (positive) information costs. Of course, the value of information will be strictly positive only if the player will actually use it. That is, there must be some states of Nature that the player can distinguish only in case he learns and in which he will play distinct strategies resulting in distinct outcomes.

In the case of private information acquisition information does not only have an informational value (as in the case of secret information acquisition) but it also has some strategic value, as it can affect other players' behaviour. In a sequential equilibrium it will often affect other players' behaviour, while in Nash equilibria other players may (pretend to) ignore the fact that some player has acquired information. As we have seen, however, such ignorance may very well be incredible. We have also seen that the strategic value of obtaining a piece of information may be strictly negative and may more than off-set the positive informational value: information may hurt. However, the information privately acquired in a sequential equilibrium never hurts.

We saw that costly information that is secretly acquired must be used. Hence, equilibria in games with exogenous information structure where private information is not used is not robust to endogenization of the information structure. Examples of such cases are pooling equilibria in signalling games and bunch-bidding in auctions. However, if information is acquired privately it is possible that costly information is acquired but not used.

Consider the example from Fig. 4. Nature chooses with equal probability between game I and II. Again, only player 1 can learn Nature's move at some small positive cost c. If given the move by player 2, player 1 has to guess Nature's move. If he guesses correctly he is rewarded but if he guesses wrong he is severely punished. In case player 1 does not know it is optimal to admit that and choose option III.

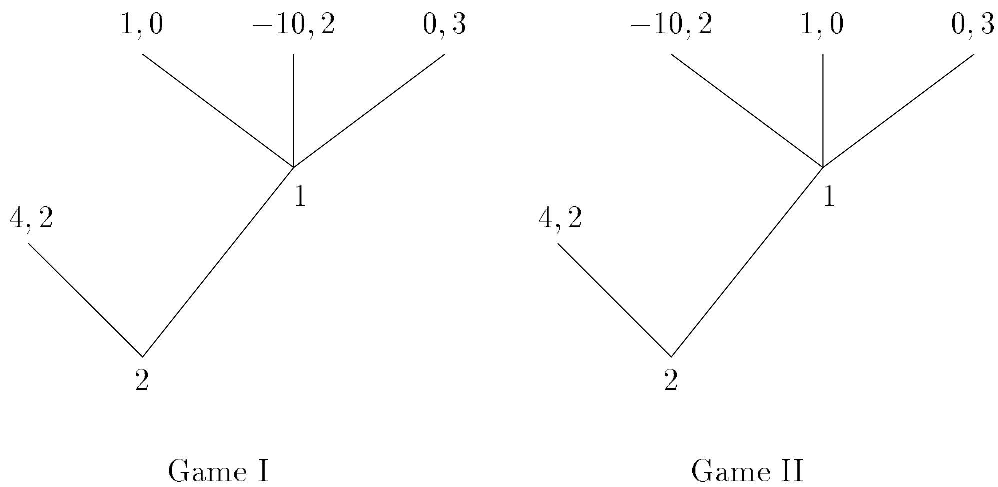

Fig. 4

It is clear that the unique sequential equilibrium of the private information acquisition game is for player 1 to learn Nature's move and guess right if given the move and for player 2 to opt out. Information is not really used as the outcome will be (4,2) independent of whether nature chooses I or II. Player 1 acquires the costly information in order to credibly threaten to guess Nature's move and thereby forcing player 2 to opt out.

In order to state the observations made in a formal way we need two definitions.

Definition 1 *We say that* all information is used *in G^P by σ if for no player i there exists a coarser information structure P_i' with respect to which either*
(i) σ_i itself is measurable, or
(ii) some other strategy σ_i' is measurable which yields the same payoff as σ_i against σ_{-i}.

Definition 2 *We say that* information hurts *player i in the equilibrium σ of G^P if all sequential equilibria of $G^{P'}$ yield player i a strictly higher payoff than σ in G^P, where P'*

is the information structure obtained from P by replacing i's information structure P_i by P^{no}.

Theorem 3 *Let G^P be a generic extensive form game with exogenous information structure $P \neq P^{no}$. Assume information costs of P are small.*

(i) Let σ be a Nash equilibrium of G^P. Then (P, σ) is a Nash equilibrium of the secret information acquisition game Γ^s if either (a) all information is used in σ or (b) information cost is zero.

(ii) Let σ be a sequential equilibrium of G^P. Then (P, σ) is a sequential equilibrium of the secret information acquisition game Γ^s if either (a) all information is used in σ or (b) information cost is zero.

(iii) Let σ be a Nash equilibrium of G^P. Then (P, σ) is a Nash equilibrium of the private information acquisition game Γ^p if information cost is zero.

(iv) Let σ be a sequential equilibrium of G^P. Then (P, σ) is a sequential equilibrium of the private information acquisition game Γ^p if information does not hurt.

Proof. (i) and (ii). Assume that any information structure other than P^{no} and P are prohibitively costly. If cost of P is zero it is clear that (P, σ) is a (sequential) equilibrium of the secret information acquisition game Γ^s. (For sequentiality define beliefs as in the game with exogenous information, that is, each player expects other players to have acquired P.)

(iii) It is clear that playing σ independent of the information acquired by others is a Nash equilibrium in the game playing phase, and that deviating in the information gathering stage is then not beneficial if information cost is zero.

(iv). If information does not hurt we can choose for each 'subgame' $G^{P'}$ (where P' is obtained from P by replacing the information structure of one player) a sequential equilibrium σ' that yields the corresponding player a lower payoff. Then it is clear that by threatening to play σ' if P' is gathered will keep all players from deviating in the information gathering phase. It is also clear that this equilibrium is sequential. □

4 Conclusion

In this paper we investigated the robustness of equilibrium results of game models with respect to the endogenization of the information structure of the players. It turned out that only when information acquisition is secret and players choose actions simultaneously in the original game, the results are fully robust as long as all information is used. When information acquisition is private, the endogenization process eliminates the equilibria of the game with an exogenous information structure where information hurts. On the other hand, when players move sequentially (and information acquisition is secret), the endogenization may generate additional sequential equilibria. In the latter case endogenous information acquisition may explain, what seemed to be irrational behavior in the game with an exogenous information structure, as rational equilibrium outcomes. The fact that we need *secret* information acquisition to obtain robustness is somewhat surprising, since in that case the information structure will not become common knowledge, while in the case of private information acquisition and games with an exogenous information structure the latter is always common knowledge.

Our results once again demonstrate the important difference between secret and private information acquisition.[7] When information acquisition is secret, information is acquired for informational purposes only. That is, it is acquired because it allows the person who possesses it to make better decisions. When information acquisition is observed, however, it may be acquired (or not) for strategic reasons: By committing to have (or not have) some piece of information the actions of other players can be influenced in a way that is favorable to the first player. Since the difference between secret and private information acquisition is so important, the choice between the two should be determined by which resembles reality best, and not so much by analytical convenience.

We have considered information acquisition as refining one's partition of the state space. In the literature private information is sometimes modelled by players receiving an imprecise signal about the true state. Also in this case one can endogenize the information structure by having players decide on the precision of information they want. We conjecture that our results also hold in this case: When information acquisition is secret and play is simultaneous, there will be a one-to-one correspondence between the results of the exogenous and endogenous information models. With sequential moves (and secret information) additional equilibria will appear, while in the case of private information acquisition some equilibria may disappear. We can even use the same examples to prove the latter statements. Just interpret no learning as information with precision zero, and full learning as information of precision 1 (or infinite). For the case where precisions can be chosen from a continuum, one would have to look further to come up with some examples, but we conjecture that there is no fundamental difference and that such examples can be constructed.

The results obtained in this paper should be useful to come up with more applications of the type discussed in the introduction where information acquisition may lead to new insights and surprising results. Assuming that information acquistion is costly and is not observed by the other players may lead to such conclusions in games with a sequential structure. A potential application would be to compare Dutch and English auctions.

[7] This point was already made, in the context of Cournot competition, by Hauk and Hurkens (1997) and Hurkens and Vulkan (1997) who focussed on the incentives to gather information.

The Dutch auction is basically a simultaneous move game, whereas the English auction is sequential as one observes all the bids made. Hence in the English auction there could be some revision of beliefs about the informedness of competing bidders which cannot occur in the Dutch auction. Future research in this direction could prove to be interesting.

5 References

Cremer, J. and F. Khalil (1992). "Gathering Information Before Signing a Contract," *American Economic Review,* **82**, 566-578.

Cremer, J. and F. Khalil (1994). "Gathering Information Before the Contract is Offered: The Case with Two States of Nature," *European Economic Review,* **38**, 675-682.

Cremer, J., F. Khalil, and J.-C. Rochet (1998). "Strategic Information Gathering Before a Contract is Offered," *Journal of Economic Theory,* **81**, 163-200.

van Damme, E.E.C. (1989). "Stable Equilibria and Forward Induction," *Journal of Economic Theory,* **48**, 476-496.

Grossman, S. and J. Stiglitz (1980). "On the Impossibility of Informationally Efficient Markets," *American Economic Review,* **70**, 393-408.

Hauk, E. and S. Hurkens (1997). "Secret Information Acquisition in Cournot Markets," UPF Economics Working Paper 233.

Hurkens, S. and N. Vulkan, (1996). "Information Acquisition and Entry," UPF Economics Working Paper 155.

Hurkens, S. and N. Vulkan, (1997). "Free Entry Does Not Imply Zero Profits," UPF Economics Working Paper 268.

Hwang, H. (1993). "Optimal Information Acquisition for Heterogeneous Duopoly Firms," *Journal of Economic Theory*, **59**, 385-402.

Hwang, H. (1995). "Information Acquisition and Relative Efficiency of Competitive, Oligopoly, and Monopoly Markets," *International Economic Review*, **36**, 325-340.

Kessler, A. (1998). "The Value of Ignorance," *RAND Journal of Economics,* **29**, 339-354.

Lee, T.K. (1982). "Resource Information Policy and Federal Resource Leasing," *Bell Journal of Economics,* **13**, 561-567.

Li, L., R.D. McKelvey, and T. Page, (1987). "Optimal Research for Cournot Oligopolists," *Journal of Economic Theory*, **42**, 140-166.

Matthews, S. (1984). "Information Acquisition in Discriminatory Auctions," in: *Bayesian Models in Economic Theory*, eds. M. Boyer and R.E. Kihlstrom. Elsevier.

Milgrom, R. (1981). "Rational Expectations, Information Acquisition and Competitive Bidding," *Econometrica*, **49**, 921-944.

Neyman, A. (1991). "The Positive Value of Information," *Games and Economic Behavior*, **3**, 350-355.

Perea y Monsuwé, A. (1997). "Buying Information in Extensive Form Games," mimeo.

Persico, N. (2000). "Information Acquisition in Auctions," Forthcoming in *Econometrica.*

Ponssard, J.-P., (1979). "The Strategic Role of Information on the Demand Function in an Oligopolistic Market," *Management Science,* **25**, 243-250.

Vives, X. (1988). "Aggregation of Information in Large Cournot Markets," *Econometrica*, **56**, 851-876.