NBER WORKING PAPERS SERIES

SORTING OUT THE DIFFERENCES BETWEEN SIGNALING AND SCREENING MODELS

Joseph Stiglitz

Andrew Weiss

Technical Working Paper No. 93

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 1990

NBER Technical Working Paper #93
November 1990

SORTING OUT THE DIFFERENCES BETWEEN SIGNALING AND SCREENING MODELS

ABSTRACT

In this paper we analyze games in which there is trade between informed and uninformed players. The informed know the value of the trade (for instance, the value of their productivity in a labor market example); the uninformed only know the distribution of attributes among the informed. The informed choose actions (education levels in the Spence model); the uninformed choose prices (wages of interest rates). We refer to games in which the informed move first as signaling games -- they choose actions to signal their type. Games when the uninformed move first are referred to as screening games. We show that in sequential equilibria of screening games some contracts can generate positive profits and others negative profits, while in signaling games all contracts break even. However, if the indifference curves of the informed agents satisfy what roughly would amount to a single crossing property in two dimensions, and some technical conditions hold, then all contacts in the screening game break even, and the set of outcomes of the screening game is a subset of the outcomes of the corresponding signaling game.

In the postscript we take a broad view of the strengths and weakness of the approach taken in this and other papers to problems of asymmetric information, and present recommendations for how future research should proceed in this field.

Joseph Stiglitz
Stanford University
Palo Alto, CA 94305

Andrew Weiss
Boston University
Boston, MA 02215

Sorting Out the Differences Between Signaling and Screening Models*

Joseph Stiglitz

Stanford University

Palo Alto, CA 94305

Andrew Weiss

Boston University

Boston, MA 02215

In recent years a good deal of interest has focused on markets with asymmetric information in which some of the participants have information that the other participants seek to acquire. Often information can be inferred from the actions taken by the informed participants. The uninformed may try to induce the informed to take actions that convey information and the choices of the informed are influenced by the information conveyed by those choices. Of course, actions also directly affect the payoffs of both the informed and uninformed. In general the actions taken and the information transmitted is likely to be sensitive to details of the economic environment. We shall focus on one aspect of the economic environment: whether informed agents move before or after uninformed agents.[1]

In models of markets with asymmetric information the equilibrium outcomes are sensitive to assumptions about how participants react to previous moves. It seems reasonable to assume that individuals react optimally; however, the definition of an optimal reaction is likely to depend on the economic context of the problem being analyzed. We shall specify below what we mean by optimal

* This paper was first written in June 1981. We have revised it slightly for this volume to take into account recent developments in the literature. We were surprised to find that the points we made at that time are not totally irrelevant to the ongoing debate on these issues. The postscript and various footnotes comment briefly on some issues that have arisen since we wrote this paper.

1 Throughout this paper we restrict our analysis to competitive environments. Clearly, equilibrium outcomes in non-competitive environments will differ markedly from those in competitive environments.

reactions: clearly optimal reactions can depend on the beliefs of the people going second concerning who made a particular move.

In much applied research on markets with asymmetric information the informed move first. These models include the Stiglitz [1982] and Bhattacharya [1980] models in which (informed) firms issue dividends that convey information about the true profitability of the company to investors who are uninformed; the Milgrom and Roberts[1982] and Salop [1979] models in which firms know their cost functions and choose a price which signals their production costs to potential entrants, who are uninformed; and the Weiss [1983] model of education in which individuals, who know their own abilities, choose a level of schooling which signals their productivity. All of these models share the characteristic that the informed participants move first, choosing a price, education level or dividend policy, and the uninformed then respond. The actions that the informed agents take may or may not fully reveal their private information. Typically these models generate a multiplicity of equilibria including some in which all the informed choose the same action (pooling equilibria) and some in which they each choose different actions (separating equilibria). Consequently models of this sort can be used to explain why individuals go to school, even if schooling is unproductive, or why firms pay dividends, despite the adverse tax effects of that practice.[2]

On the other hand, in some models unreasonable equilibria emerge that seem due to peculiarities of the model or the definition of equilibrium employed rather than the underlying structure of the markets. For example, suppose that firms believe that if anyone chooses other than 8 years of education that person has zero productivity, then equilibrium will be characterized by all individuals choosing 8 years of education. These beliefs

2 The literature arguing that dividends are paid because they provide a signal concerning the firm's net worth is, however, not completely persuasive. Presumably, buying back shares would provide an equally effective signal, at much lower cost. On the other hand, we show below that there may exist signalling equilibria which are far from Pareto efficient. Perhaps the dividend signalling equilibrium is a dramatic example of this.

are unreasonable. Or suppose investors believe that any company choosing a dividend pay out rate below 5% of net asset value is in imminent danger of bankruptcy and value the stock accordingly, then one equilibrium is characterized by all solvent firms paying a 5% dividend.

Parallel to the treatment of markets in which informed agents move first—make choices to which the uninformed respond—have been analyses of markets in which the uninformed participants move first. Early treatments of this problem were by Stiglitz [1975] Riley [1977,1979], Rothschild and Stiglitz [1976] and Wilson [1977].[3] In the Stiglitz and Riley papers uninformed firms offer wage contracts—a wage conditional on an education level—and informed individuals react to those wage contracts by choosing the education level that maximizes their utility. In the Rothschild-Stiglitz-Wilson papers uninformed insurance companies offer contracts and customers choose their most desirable contract given their probability of an accident and their risk preferences.[4]

We refer to models in which the informed move first as signaling models—the more desirable informed agents signal who they are.[5] We refer to models in which the uninformed move first as screening models—contracts are designed to screen the more desirable agents from the less desirable ones. One quality shared by signaling and screening models is that they generate surprising and often counterintuitive results. In signaling models there are often multiple

3 The order of moves in the Spence model [1973,1974] is somewhat ambiguous. Spence [1976] interprets Spence [1973]as the informed agents choosing education levels before firms make wage offers. One could also view the original Spence model as a simultaneous move game.

4 Bhattacharya [1980], Weiss[1980], Gausch and Weiss [1980,1982], and Salop and Salop [1976], Lazear and Rosen [1981] and Nalebuff and Stiglitz [1983] among others have modeled labor markets where uninformed firms first offer wage contracts and individuals then apply to the firm offering the most advantages contract given their characteristics.

5 Since we first circulated this paper this terminology has become commonplace. Spence [1976] refers to a model in which the informed move first as a passive response model.

(optimal reaction) equilibria, some of which seem implausible; in screening models a pure strategy equilibrium often does not exist.[6] That is, there is no set of contracts offered by all the uninformed agents that would not induce at least one uninformed participant to offer a contract different from the one assigned him. Thus it would appear that the set of outcomes when the uninformed move first is a subset of the set of outcomes when the informed move first (a claim along these lines was made in Spence [1976] in the context of active and passive responses to signals.) This conjecture turns out not to be strictly correct. In particular when the uninformed move first (sequential) equilibrium outcomes may be characterized by some contracts generating positive profits and others generating losses for the uninformed. These outcomes cannot arise when the informed move first. Only when the parameter values of the problem are such as that these contracts are not offered in equilibrium, are we able to prove an inclusion relationship.

In the screening literature the equilibrium notions customarily used implicitly impose optimal reactions by the informed agents to any contracts that are offered (whether in equilibrium or not). The signalling literature faces a more difficult problem: informed agents (who move first) must make an inference about how the uninformed agents will respond to any action (including any out-of-equilibrium action) they take. What are the "optimal" responses of uninformed agents depends on the inferences they draw, and it is not always clear what those inferences will, or "should", be when there are heterogeneous agents. This is particularly true for "out of equilibrium" actions; for the theory predicts that no rational agent will take those actions. Consequently, in the signaling literature the optimal reaction assumption is not always imposed (for example, it is not present in Spence's signaling models.)

In this paper we define optimal reaction equilibria for both screening and signaling models. For screening models the optimal reaction equilibria

6 See Rothschild-Stiglitz [1976], Guasch and Weiss [1980, 1982], Bhattacharya [1980] and Riley [1979] for examples of models in which equilibria fail to exist.

are subgame perfect or Stackelberg equilibria with the uninformed as leaders and the informed as followers. (They can also be described as two stage games in which the uninformed move first.)

For signaling models, optimal reaction equilibria are roughly equivalent to sequential equilibria appropriately modified to allow for continuous action spaces. Following Kreps-Wilson [1982] we implicitly assume that all uninformed agents have the same beliefs and that they react optimally given these beliefs, (and the strategies of all agents) to the observed actions. Beliefs assign to each action a probability distribution of agents taking it, and have the following properties:

(a) the equilibrium combination of strategies and beliefs cannot contradict one another: if the equilibrium strategy combination calls for only player i to choose action a then uninformed agents seeing action a must believe that it was chosen by player i;

(b) if out of equilibrium actions were to be observed, uninformed agents could not believe that those actions were taken by agents that are not present in the economy or by agents for whom they are not feasible;

(c) beliefs cannot be contradicted by the observed distribution of actions. For instance, suppose there are two equal sized groups I_1 and I_2 of informed individuals. Members of groups I_1 have single feasible action a_1 (the only inference consistent with the observation that precisely half the informed chose a_1 is that actions other than a_1 were taken only by members of group I_2).

As we shall see in the examples in section 2, assumption c is quite strong. However it is necessary if beliefs are to satisfy Bayes Rule. These restrictions on beliefs motivate the restrictions we place on strategies when we formally define optimal reaction equilibria. As one might expect, in many models there will be several optimal reaction equilibria. On the other hand, the optimal reaction restrictiona eliminate some of the least reasonable Nash equilibria in the same way they are eliminated by imposing sub-game perfection or sequentiality.[7]

The principal result of this paper is that if we restrict ourselves to economies in which the only contracts realized in equilibrium are ones that break even, than the set of outcomes of the optimal reaction equilibria when the uninformed move first are a subset of the optimal reaction equilibria when the informed move first. But the restrictions needed to eliminate positive profit contracts are surprisingly strong.

This result can best be understood by considering reasonable reactions to out-of-equilibrium moves in each game. When the uninformed move first, the optimal reactions of the informed to out-of-equilibrium moves are dictated solely by the preferences of the informed agents. When the informed move first, however, the optimal reactions of the uninformed to out-of-equilibrium moves depend on their beliefs about which agent(s) took those moves as well as on their preferences. Hence, there is more leeway for "bad" reactions to out-of-equilibrium moves when the informed move first than when the uninformed move first.

In an optimal reaction equilibrium when the uninformed move first unreasonable reactions are precluded. The informed always choose their most desirable contracts from the set of contracts being offered; and the uninformed know this. Hence the uninformed have no uncertainty about the matching of informed agents to actions in response to any set of contracts. Profit maximizing behavior by the informed precisely determines the actions that would be chosen in response to all price schedules, including those not offered in equilibrium.

When the informed move first it is possible for the uninformed to believe that if an out-of-equilibrium action were chosen, that it was chosen by the

7 In many game theoretic formulations of general signaling games stronger restrictions on beliefs are imposed such as the Cho and Kreps intuitive criterion, or Divinity or Universal Divinity in Banks and Sobel. There is some controversy over whether these stronger restrictions are not too strong. We have chosen to make weak assumptions about beliefs and allow the reader to draw upon the particular features of the market(s) that interest him to justify stronger restrictions on beliefs.

least desirable informed agent. These pessimistic beliefs could deter informed agents from departing from their assigned actions. Consequently by allowing for pessimistic beliefs, outcomes that are precluded in the optimal reaction equilibria when the uninformed move first may be sustained as equilibria when he informed move first. Since pessimistic beliefs hurt informed agents taking the action to which those beliefs apply, we consider those beliefs as punishing the agents taking the associated out-of-equilibrium action, and thus enforcing the equilibrium.

I. A GENERAL MODEL

In this section we describe the general class of markets with which we are concerned. Because the structure of the model depends on the order in which moves occur, we postpone our discussion of the strategies of players in each game (informed moving first, uninformed moving first and simultaneous moves) until after we have described the preferences and available actions and information of the participants. We will allow agents to only pursue pure strategies.

There are finite sets I and K of informed and uninformed agents respectively. Sets I and K each have at least 2 members. Informed and uninformed agents trade with one another, and the terms of trade can be predicated on the action taken by an informed agent.

Each informed agent i chooses an action $a \in A_i$, where all A_i are compact sets in R^n. We define A as $\cup A_i$. Action a has cost $c(a,\{i\}) \in R^1$ for agent i. We refer to all informed agents with the same feasible set A_i and same $c(a,\{i\})$ functions as being the same type. Each informed agent makes one transaction (e.g. chooses a level of education and works for a single firm.)

There are constant returns to scale in transactions for uninformed agents, so the number of trades an agent makes does not directly affect his net payoff per transaction. Each uninformed agent k chooses a price p for each action $a \in A$.[8] The price p is the monetary transfer from the uninformed

8 This is a restriction on the strategy space of uninformed agents. For instance, it rules out strategies in which the uninformed agent fixes the ratio of the numbers of trades he is willing to engage in at different prices.

to the informed, and may be negative. Uninformed agents are uninformed only about the identity of the informed. If an action is feasible for more than one informed agent, uninformed agents cannot discriminate among agents choosing that action.

We allow uninformed agents to be either buyers or sellers—so that prices refer either to the prices they pay as buyers or the prices they receive as sellers. (When the uninformed are sellers p is typically negative.) In sorting models of the education-employment market, uninformed firms are buyers of labor services. The action is an education level chosen by individuals, and the price is the wage that a firm offers to pay workers with a given education level. The reader will find it helpful to keep the education example in mind throughout most of this paper. In the Rothschild-Stiglitz and Wilson models of the insurance market, uninformed insurance companies are sellers of insurance. The action is the amount of insurance customers demand, and the price is the cost of insurance for a customer demanding a given amount of coverage, and $-c(a,\{i\})$ is the value individual i places on "a" units of insurance coverage.

The expected value to an uninformed agent from a trade with an informed agent randomly selected (with equal probability) from the set J choosing action a is $\phi(a,J)$. In the education example $\phi(a,\{i\})$ is the expected value of the labor input of individual i with "a" years of education." The expected payoff of this trade for an uninformed agent offering price (wage) p for that action is $\phi(a,J)-p$. The uninformed are buyers: p is the wage and $\phi(a,J)-p$ is the firm's expected profit per worker hired with education level "a" and paid wage p. When the uninformed are sellers, $\phi(a,J)$ is generally negative and refers to the cost of providing the good (or service) to a buyer that is randomly selected from set J. However, we would again emphasize that the paper can be most easily followed by keeping in mind the education example in which the uninformed are firms hiring workers, the actions are education levels chosen by workers, and the prices are the wages paid by firms to workers.

To ensure that there always exists a "best" action or set of actions for informed agents, so that payoffs are defined for all combinations of prices, we impose the technical restriction that any price schedule offered by an uninformed agent must be upper semicontinuous.

The preferences of agents are to maximize their expected payoffs.[9] This implies that informed agents always trade with the uninformed agent offering the highest price for their selected actions. The payoff for informed agent i choosing action $\hat{a}$ when p is the highest price offered for action $\hat{a}$ is $p - c(\hat{a}, i)$. In the education employment example this is the worker's wage net of his cost of education. We adopt the following tie-breaking rules: if k uninformed agents are offering the maximum price (wage) for action (education) $\hat{a}$, an informed agent choosing action $\hat{a}$ trades with each of them with probability $1/k$. If the net payoffs for contracts offered in equilibrium are identical at two or more different actions, we assume that informed agents choose the action at which the profits of the uninformed are highest. Finally we assume that an agent participates in the market if and only if the expected payoff from participation is greater than or equal to zero.

There is common knowledge about the parameters, and distribution of agents in the economy. In particular, all agents know the elements of I, each A_i, and for $\forall a$ the values of $c(a,\{i\})$, and $\phi(a,\{i\})$.

CASE 1: UNINFORMED AGENTS MOVE FIRST

In this case the uninformed choose a price schedule, the informed then choose actions. Finally each informed agent automatically trades with the uninformed agent(s) whose contract yields the highest payoff to that informed agent (with ties broken as above.) We shall not allow the uninformed agents to predicate the price schedule they offer upon the subsequently observed

9 These are simplifying assumptions. Our results are valid for a more general class of preferences $v(a, i, p)$ as would be required in the insurance example.

distribution of realized actions of the informed agents.[10] Thus the strategy of uninformed agent k is the price schedule $P_k: A \rightarrow R^1$. This specifies the price the k^{th} agent offers to any informed agent choosing action a. The upper envelope of these price schedules is an upper semicontinuous function denoted by P; that is, $\forall a$, $P(a) = \max_k P_k(a)$. Since the informed agents undertaking action a always trade with the uninformed agent offering the highest price for the action, trades will only occur along the price locus P.

The strategies of informed agents are potentially more complicated: they observe all the prices offered before choosing an action (though only the upper envelope of those price schedules is relevant for their payoffs.) Let Π denote the set of feasible combinations of price schedules and π be an element of Π. Then a strategy combination for informed agents is described by a function f: $I \times \Pi \rightarrow A$; $f(i,\pi)$ describes the action chosen by agent i when the combination of price schedules π obtains. $f^{-1}(a,\pi)$ denotes the set of informed agents choosing action a under f when π is the combination of price schedules being offered. For any set of price schedules π and strategy combination of informed agents, the set of actions a for which $f^{-1}(a,\pi)$ is non-empty is denoted by $\tilde{A}$.

Definition 1. *A Nash equilibrium when the uninformed move first is a combination of strategies* $(P_1^*..P_m^*,f^*)$ *such that, given the strategies of all other players, no uninformed agent k could increase his expected payoff by offering a price schedule* $\hat{P}_k \neq P_k^*$, *and no informed agent i could increase his payoff by choosing an action* $a \neq f^*(i,\pi^*)$, *where* $\pi^* = (P_1^*...P_k^*)$.

This definition of equilibrium places no restrictions on the reactions of informed agents to combinations of price schedules other than π^*.

10 If the uninformed were able to make the price they offer for action a be a function of the distribution of actions taken by the informed, the distinction between the informed moving first and the uninformed moving first would be blurred. By precluding contingent contracts of that form we preserve the distinction between the informed moving first and the uninformed moving first.

In the context of the education example, suppose $\forall_i$: $\theta(8,i) - c(8,i) > 0$. One Nash equilibrium would be for (uninformed) firms to offer a wage equal to the average productivity of a randomly selected worker if all workers were to have 8 years of schooling, and a zero wage to any worker choosing other than 8 years of schooling and for all (informed) individuals to choose 8 years of education regardless of the wage offers of firms. Note that any individual choosing other then 8 years of education is worse off than if he had chosen 8 years of education, while no firm is better off by offering a positive wage for education levels other than 8 since no individual chooses education levels different from 8 years.

Clearly this equilibrium is unreasonable. Analyses of models of asymmetric information when the uninformed move first have eliminated equilibria of this sort by implicitly or explicitly imposing an optimal reaction assumption. Optimal reactions such as those generated in a Stackelberg or subgame perfect equilibrium of this game assume that whatever price schedules are offered, each informed agent must react by choosing the action that maximizes his expected payoff. These reactions are anticipated by the uninformed agents before they offer contracts; hence, they determine responses to possible out-of-equilibrium moves. The reasoning behind this restriction is that it is reasonable to expect the informed agents to choose actions which yield the maximum payoff for any price schedule.

Definition 2. *$(P_1^*, \ldots, P_n^*, f^*)$ is an optimal reaction equilibrium when the uninformed move first (ORUF) if $(P_1^*, \ldots, P_n^*, f^*)$ satisfy the conditions for a Nash equilibrium and $\forall_i \in I$, $\pi \in \Pi$, $f^*(i,\pi) = a \in \arg\max_{a \in A_i} (P^*(a) - c(a,i))$. If $\arg\max_{a \in A_i} (\cdot)$ has more than one member, agent i chooses the action in that set which maximizes $\theta(a,i) - P^*(a)$. If several of those actions maximize $\theta(a,i) - P^*(a)$ they choose each with equal probability.*

The ORUF coincide with the sub-game perfect equilibria. The additional restrictions imposed by ORUF are particularly compelling since they only eliminate equilibria which use dominated strategies.

CASE 2. INFORMED AGENTS MOVE FIRST

In this case the informed first choose actions. The uninformed then offer price schedules. Finally, each informed agent automatically trades with the uninformed agent whose contract gives the highest payoff to the informed. For example, individuals first go to school, firms then offer wages conditional on years of education, and, finally, each individual goes to work for the firm offering the highest wage given the selected level of education. Since the informed choose their actions before the uninformed choose price schedules, the actions of the informed agents cannot depend on the price schedules of the uninformed. For this game we shall let ψ denote a strategy combination for the informed agents; $\psi: I \rightarrow A$. Thus $\psi(I)$ describes an action chosen by agent i, and $\psi^{-1}(a)$ denotes the set of agents choosing action a. The set of actions for which $\psi(I)$ is non-empty is denoted by $\tilde{A}$. (The use of the same notation as in case 1 eases the exposition.) On the other hand the uninformed agents choose price schedules after having observed the distribution of realized actions of the informed agents. Let T denote the set of observable distributions of the actions of the informed agents, $t \in T$, is a particular observed distribution of actions.

Turning now to the strategies of the uninformed agents, let ρ_k denote the price offered by agent k to any informed agent choosing action "a" when the distribution of actions is t, so that ρ_k: $A \times T \rightarrow R^1$. That is, for each $t \in T$, agent k may choose a different upper semi-continuous price schedule. Note that a price may be offered for actions which were not chosen; those prices are irrelevant for the equilibrium payoffs but do affect whether a strategy combination is an equilibrium.

Definition 3. *A Nash equilibrium when the informed move first is a combination of strategies* $\{\rho_1^* \ldots \rho_m^*, \psi^*\}$ *such that, given the strategies of all other agents, no uninformed agent k could increase his expected payoff by offering a price schedule* $\hat{\rho}_k \neq \rho_k$ *and no informed agent i could increase his payoff by choosing action* $\hat{a} \in A_i$ *such that* $a \neq \psi(i)$.

As in the previous game the Nash definition allows for unreasonable equilibria. We shall now apply the logic behind the Kreps- Wilson notion of a sequential equilibrium to define an optimal reaction equilibrium for this game.[11] We consider reasonable restrictions on $\rho(.)$, the maximum price offered by the uninformed agents in response to different actions. The motivation for the restrictions we impose is:

a) All the uninformed agents have the same beliefs about the probability distribution of agents choosing an out-of-equilibrium action, if such an action were observed;

b) these beliefs are consistent with the feasible action spaces of each informed agent and the observed distribution of actions, t.

Definition 4. *An optimal reaction equilibrium when the informed agents move first (ORIF) is a Nash equilibrium with the additional property that* $\forall a \in \bar{A},\ t \in T,\ \rho(a,t) \geq \min_{\tilde{I}_{a,t} \subset I} \theta(a, \tilde{I}_{a,t})$ *where* $\tilde{I}_{a,t}$ *is any non-empty subset of the informed agents for whom action a is feasible, and* $\tilde{I}_{a,t}$ *choosing a is consistent with having observed t.*

In the context of the education example, the definition of ORIF precludes the highest wage offers made for out-of-equilibrium education levels being below the value of the lowest labor input that any worker could possibly achieve at that education level.

In the special case where *A* is discrete, the outcomes of the set of ORIF coincide with the outcomes of sequential equilibria. (Sequential equilibrium

11 When the informed move first there are no proper subgames. Consequently subgame perfection does not reduce the set of Nash equilibria.

is not defined when A is continuous.) Our definition avoids both the additional notation required by beliefs in Kreps–Wilson and the complexity inherent in applying their concept to games with continuous action spaces.[12]

II. SOME EXAMPLES

We can most readily show how the ordering in which players move affects the trades realized as optimal reaction, or Nash, equilibria through a series of simple examples. These examples are intended to serve as pedagogical tools to illustrate some problems that have arisen in analyses of markets with asymmetric information. Since the point of these examples is purely illustrative, we shall not describe the outcomes for every definition of equilibrium and ordering of moves. To ease the notation and shorten the exposition, when no confusion would result, we shall use the notation $p_k(a)$ and $p(a)$ to describe price offers that are independent of the observed combinations of actions. (When the notation $p(a)$ is used for a game in which the informed move first, the reader should assume that all uninformed agents are offering the same price schedule, and the strategies of the uninformed agents are such that prices are independent of the observed distribution of actions.) This section can be skipped without loss of continuity. Interested readers may wish to glance at example 4 where positive profit contracts are outcomes of ORUF equilibria.

Example 1.

The first example illustrates cases in which there are Nash equilibrium of both the signaling and screening models that are clearly unreasonable. There are two types of informed agents I_1 and I_2. A type I_1 agent can only

12 The restrictions we have imposed on the strategies of the uninformed in the ORIF equilibria are quite weak, yet, as we shall see, they are not sufficient to generate the subset relationship ORUF $\subset$ ORIF conjectured by Spence. Cho and Kreps impose stronger restrictions in a somewhat different signaling game and show that given their restrictions the outcome when the informed move first coincides with the outcome when the uninformed move first (if the latter exists).

choose action a_1. A type I_2 agent can choose either action a_1 or a_2. All actions are costless. The expected value to an uninformed agent of trading with a randomly selected type I agent choosing action j is denoted by θ_{ji}. The expected value to an uninformed agent of trading with a randomly selected individual choosing action a_1 when that action is chosen by the entire population is $\bar{\theta}_1$.

Suppose $0 < \theta_{12} < \bar{\theta}_1 < \theta_{11} < \theta_{22}$. Thus the uninformed value trade with a type 2 informed agent choosing action a_2 more highly than with a type 1 informed agent choosing action a_1. They also value a trade with a type 2 agent choosing action a_2 more than if that agent chose action a_1. One characterization of the moves in a Nash equilibrium, regardless of who moves first, is for both informed types to choose action a_1, and for the uninformed to offer prices $p(a_1) = \bar{\theta}_1$ and $p(a_2) = 0$. The strategies of the agents in each game are to make these moves regardless of any observation of the moves of other agents. An uninformed agent offering $P_k(a_1) > \bar{\theta}_1$ would generate losses, $P_k(a_1) < \bar{\theta}_1$ would not attract any informed agents, $P_k(a_2) \neq 0$ would also not attract any informed agents —recall that the definition of Nash equilibrium holds the strategies of all other agents fixed, and all informed agents are choosing action a_1 regardless of the observed prices. Similarly no informed agent gains from deviating from action a_1 given the equilibrium contracts offered by the uninformed. Although the strategy choices of the uninformed satisfy the criteria for a Nash equilibrium they seem unreasonable. Only type 2 agents can choose action a_2; one would imagine that in equilibrium they would choose that action and be appropriately rewarded. A more reasonable combination of moves that is also sustained as a Nash equilibrium, regardless of which agents move first, is for the type 1 informed agents to choose action 1, type 2 to choose action 2 (again these choices are made independently of observed prices) and for uninformed agents to offer prices $p(a_1) = \theta_{11}$, $p(a_2) = \theta_{22}$.

When the uninformed move first, there is also a positive profit (albeit also unreasonable) Nash equilibrium for this example.

$$\pi^* \begin{cases} \forall k,\ P_k^*(a_1) = \theta_{11} \\ \forall k,\ P_k^*\ (a_2) = \bar{\theta} \quad \text{where } \theta_{11} < \bar{\theta} < \theta_{22} \end{cases}$$

$$f^*(I_1,\pi^*) = a_1$$

$$f^*(I_2,\pi^*) = a_2$$

$$f^*(I_2,\pi\ ) = a_1,\ \text{for } \pi \neq \pi^*$$

where I_1 denotes a type 1 agent, and I_2 denotes a type 2 agent. In this Nash equilibrium the contracts offered to agents choosing action a_1 break even, the contracts offered to agents choosing action a_2 make positive profits. Any contract offering a higher price than $\bar{\theta}$ for action a_2 would precipitate a move by the type 2 agents <u>away</u> from a_2 to a_1. This irrational response by the informed agents would decrease the expected payoff of the uninformed agent deviating from that equilibrium. (Recall that irrational responses to out-of-equilibrium moves are permitted by the definition of a Nash Equilibrium.)

Let us now consider the case where the informed move first and the uninformed react optimally (ORIF). The only ORIF equilibrium is

All agents of type 1 choose a_1

All agents of type 2 choose a_2

$$\forall k \in K,\ t \in T,\ \begin{cases} \rho_k(a_1,t) = \theta_{11} \\ \rho_k(a_2,t) = \theta_{22} \end{cases}$$

Both types choosing a_1 is not an equilibrium because the informed know that profit maximizing behavior by the uninformed will cause them to offer $\rho(a_2,t) = \theta_{22}$ if action a_2 is observed (this is the optimal reaction which is consistent with profit maximizing behavior by the uninformed agents).

Similarly if the uninformed agents move first, the only optimal reaction equilibrium is the separating one, $P(a_1) = \theta_{11}$, $P(a_2) = \theta_{22}$. An uninformed agent knows that if uninformed agents offer the pooling contract $P(a_1) = \bar{\theta}_1$, then a contract $P_k(a_2) = \theta_{22}-\epsilon$, with $0 < \epsilon < \theta_{22} - \bar{\theta}_1$ will attract only type 2 agents and earn positive profits. Thus, in this example, the realized actions and payoffs in the optimal reaction equilibrium are the same regardless of who moves first.

Example 2.

We shall now modify example 1 to show that both ORIF and ORUF equilibria may fail to maximize output. We now assume action a_2 is available to type 1 as well as to type 2, and let the expected value of a trade with type 1 choosing a_2 be θ_{21}.

Suppose that

$$0 < \theta_{21} < \theta_{12} < \bar{\theta}_1 < \bar{\theta}_2 < \theta_{11} < \theta_{22}.$$

If the informed agents move first there are several combinations of moves characterizing optimal reaction equilibria (ORIF) in this example.

(i) Both types choose a_1 and $\forall t \in T$,

$$\rho(a_1,t) = \bar{\theta}_1$$

$$\rho(a_2,t) = \theta_{21}$$

Implicitly, the uninformed expect that if action 2 is chosen it is chosen by type 1.

(ii) Both types choose action a_2, and $\forall t \epsilon T$,

$$\rho(a_2,t) = \bar{\theta}_2$$

$$\rho(a_1,t) = \theta_{12}$$

In this case the uninformed expect that if action a_1 is chosen, it is chosen by type 2. Both of these pooling equilibria are sustained by pessimistic beliefs by the uninformed: out-of-equilibrium actions are only taken by type 1.

(iii) If we were to change our assumption on how informed agents choose actions in cases of indifference there is a third category of ORIF equilibria.

Those equilibria are characterized by action a_1 and a_2 being chosen by the informed agents in proportions such that the value of a trade with a randomly selected agent choosing either action is identical. For instance, one ORIF equilibrium has all type 1 agents choosing a_2 and the proportion of type 2 agents choosing action a_2 being such that the values of a trade with a randomly selected agent choosing action a_2 is $\bar{\theta}_2 = \theta_{12}$. The uninformed then offer $P(a_1) = P(a_2) = \theta_{12}$. We have not restricted the expectations of the uninformed except that (a) given their expectations the strategies of the informed must be consistent with those expectations, and (b) the expectations must be feasible. In each of these ORIF equilibria, resources are being misallocated. Agent i does not necessarily choose the action which maximizes $\theta(a,\{i\})$.

None of the ORIF equilibria described thus far maximize output. Output is maximized if type 1 choose a_1 and type 2 choose a_2. Those choices will only emerge as an ORIF equilibrium if the self-selection constraint is violated—in particular type 1 would prefer the contract received by type 2 to their own contract. Because of our assumptions that the distribution of realized actions is observed by the uninformed before they choose price schedules, and that there is a finite number of informed agents it is possible to construct such an equilibrium. We propose these equilibria as curiosities. They depend on the action chosen by a single individual affecting the distribution of prices paid by the informed.

Suppose there are ℓ type 2 agents. Let $\hat{\theta}_2$ denote the belief the uninformed agents have about the expected value of a trade with an agent choosing action a_2 if the number of agents choosing a_2 differs from ℓ. Then if $\hat{\theta}_2 < \theta_{11}$ there is an output maximizing ORIF equilibrium:

Type 1 chooses a_1.

Type 2 chooses a_2.

For t such that ℓ agents choose a_2, $\rho(a_2,t) = \theta_{22}$ and $\rho(a_1,t) = \theta_{11}$.

For t such that $m \neq \ell$ agents choose a_2, $\rho(a_2,t) = \hat{\theta}_2$ and $\rho(a_1,t) \leq \theta_{11}$.

One motivation for these strategies is that the uninformed believe that if they observe $m \neq \ell$ agents choosing action a_2 then the proportion of type 1 agents choosing a_2 is sufficiently large that the expected productivity of an agent randomly selected from among those choosing a_2 is $\hat{\theta}_2$.

This ORIF equilibrium violates the self selection constraint usually imposed in models with asymmetric information. Type 1 agents are choosing action a_1 although, holding the contracts fixed, they would do better if they chose action a_2. However, if any type 1 agent were to switch to action a_2 the distribution of actions observed by the uninformed would change causing them to revise their beliefs in such a way as to make the type 1 agent regret having switched to action a_2.

The existence of an ORIF equilibrium that violates the self-selection constraints does not depend on the uninformed observing the entire distribution of actions. If they only observed the support of that distribution—which actions were chosen— an ORIF equilibrium could still violate the usual self-selection condition. This result can easily be seen in the context of our example by assuming that there is only one type 1 agent. The strategies of the informed are that the type 1 agents choose action a_1. The strategies of the uninformed agents are:

$$\text{If } a_1 \text{ is observed } \rho(a_2,t) = \theta_{22}$$

$$\rho(a_1,t) = \theta_{11}.$$

$$\text{If } a_1 \text{ is not observed } \rho(a_2,t) = \bar{\theta}_2.$$

Again although the type 1 agent prefers the contract offered for action a_2, if he stopped choosing a_1, then action a_1 would not be observed, and the distribution of price schedules would change accordingly.

If the uninformed move first there is only one optimal reaction equilibrium in which each contract breaks even: the uninformed agents offer $P(a_2) = \bar{\theta}_2$ and $P(a_1) < \bar{\theta}_2$, and both types of informed agents choose action a_2. Note that in this ORUF equilibrium type 1 agents are choosing the action which *minimizes* $\theta(a, I_1)$. Resources are being misallocated. Because actions are costless there cannot be an ORUF equilibrium where the two types of agents choose different actions. Thus the output maximizing actions cannot be outcomes of an ORUF equilibrium.

Thus far neither of our examples have dealt with what is often considered the quintessential problem in the analysis of market with imperfect information — nonexistence of equilibrium. To discuss that problem we need to expand the action space of the informed agents.

Example 3.

This example illustrates the Spence model of the education market. The informed agents are individuals and the action they choose is an education level. The uninformed are firms offering a wage level. The standard result that equilibrium may not exist in this market only applies to ORUF equilibria. Spence's original specification of multiple equilibria in this market can be formally justified if we consider ORIF or Nash equilibria—regardless of the ordering of moves. To be precise, suppose the set of actions available to the informed agents is the open interval $A = (0, \bar{a})$ and actions are costly, where cost to a type i individual of action a is $c_i a$ and $c_1 > c_2$. We simplify the exposition by following Spence in assuming that actions do not affect the value of a trade with a given type. The value of a transaction with type i is denoted θ_i, the expected value of a transaction involving a randomly chosen individual from the population is $\bar{\theta}$ and $\theta_1 < \bar{\theta} < \theta_2$, and $\theta_2 - c_1\bar{a} < 0$. In this example there are again many Nash equilibrium outcomes regardless of the ordering of moves. For instance all informed individuals choosing the same action $\hat{a} \in A$ and the uninformed offering prices

$$p(\hat{a}) = \bar{\theta}$$

$$\text{and for } a \neq \hat{a},\ p(a) = 0$$

is a Nash equilibrium for any $\hat{a} \in A$ (and $t \in T$ for the case when the informed move first). There is also a Nash equilibrium characterized by type 1 choosing $a = a^*$,

$0 \le a^* \le \tilde{a} - (\theta_2-\theta_1)/c_1$, type 2 choosing $\tilde{a} = \frac{\theta_2-\theta_1+c_1a^*}{c_1}$, and the uninformed choosing prices $p(\tilde{a}) = \theta_2$, $p(a^*) = \theta_1$, and $p(a) = 0$ for $a \ne (a^*,\tilde{a})$. A type 1 is no better off choosing $a \ne a^*$, a type 2 is worse off choosing $a \ne \tilde{a}$, and no uninformed participant can be made better off, given the strategies of the other participants, by offering contracts other than those specified. Let us now consider the optimal reaction equilibria when the informed agents move first (ORIF). The set of optimal reaction equilibria where both types choose the same action $\hat{a}$ (pooling equilibria) are characterized by for $\forall t \in T$, $\hat{a} \in A$, $\rho(\hat{a},t) = \bar{\theta}$

$$\text{and } a \ne \hat{a},\ \theta_1 \le \rho(a,t) \le \bar{\theta} - c_1\hat{a}$$

There is also a separating equilibria. Type 1 chooses $a^* = 0$ and type 2 $\tilde{a} = \frac{\theta_2-\theta_1}{c_1}$. In this example, the restriction imposed by the optimal reaction condition only eliminate some (not all) of the separating Nash equilibria when the informed agents move first.

On the other hand in this example there may not exist an ORUF equilibrium. First, we know there cannot be a pooling equilibrium where both types choose the same action. The proof is by contradiction; a pooling equilibrium where both types choose action $\hat{a}$ would be characterized by $p(\hat{a}) = \bar{\theta}$.

Let some uninformed agent k offer a contract $p_k(\tilde{a}) = \tilde{\theta}$, $\theta_2 > \tilde{\theta} > \bar{\theta}$ and $\tilde{a} > \hat{a}$, such that

$$c_1[\tilde{a} - \hat{a}] > \tilde{\theta} - \bar{\theta} > c_2[\tilde{a} - \hat{a}]$$

This contract attracts only type 2 individuals and makes positive profits. Notice that for this example generic nonexistence of a pooling optimal reaction equilibrium when the uninformed move first depends on the set of actions available to the informed being an open set. To break the pooling equilibrium in our example it is necessary that there exist some $\tilde{a} > \hat{a}$. If A were a closed set $[0,\bar{a}]$, and if further assumptions were made concerning θ_2

and $c_1\bar{a}$, an ORUF equilibrium could be characterized by all informed agents choosing action $\bar{a}$. (In that case a competitor cannot break the equilibrium by offering a contract contingent on an action greater than $\bar{a}$, since those actions are not feasible.)

These boundary problems could also be avoided by assuming that $\theta_2 - c_1\bar{a} < 0$. Then an individual rationality constraint precludes an ORUF pooling equilibrium at $\bar{a}$.

An optimal reaction equilibrium in which each type of informed agent chooses a different action (sorting) does not exist if, for all actions $a \epsilon A$ satisfying

$$\theta_2 - ac_1 \leq \theta_1$$

(a necessary condition for type 1 to be dissuaded from choosing the same action chosen by type 2) it is the case that

$$\theta_2 - ac_2 < \bar{\theta}.$$

In that case for an equilibrium contract to only attract type 2, it must require an action so large that if an alternate contract $P_k(0) = \bar{\theta} - \epsilon$ were offered, ϵ could be made sufficiently small that it would attract both types and make positive profits.

Example 4.

Our final example illustrates an ORUF equilibrium that is characterized by some contracts generating positive profits and others generating losses. As such it clashes with analyses which were motivated by the same considerations as our ORUF yet in which part of the definition of equilibrium was that each contract generates zero profits.

We assume there are three types of informed agents (1,2,3) and an interval of available actions $A = [1,3]$ $\theta_{ji} = 10$ for $i = j$ $\theta_{ji} = 0$ for $i \neq j$. $c(j,i) = 2|a - i|$ so that the cost of action j for agent i is twice the absolute difference between the agent and the action.

There are twice as many type 2 agents as either type 1 or type 3, (there are equal numbers of type 1 and type 3).

Finally there are 2 uninformed agents (i,j). Their strategies are

$$P_i(a_1) = 11$$

$$P_i(a_2) = 9$$

$$\text{for } a \neq (a_1, a_2),\ P_i(a) = 0$$

$$P_j(a_2) = 9$$

$$P_j(a_3) = 11$$

$$\text{for } a \neq (a_2, a_3),\ P_j(a) = 0$$

The optimal reactions of the informed agents are

type 1 chooses a = 1.

type 2 chooses a = 2.

type 3 chooses a = 3.

Each uninformed agent earns positive profits from contract $P(a_2)$, but if he raised the price he offered he would attract a different type and generate losses.

Diagrammatically, the situation is illustrated in Figure 1. The locii I_1, I_2, I_3 plot the indifference curves of types 1, 2, 3 respectively through the contracts they choose in equilibrium.

Each contract offered by the uninformed earns zero profits and resources are efficiently allocated (there is no dead-weight loss.) No uninformed agent can offer a new set of contracts that would make positive profits.

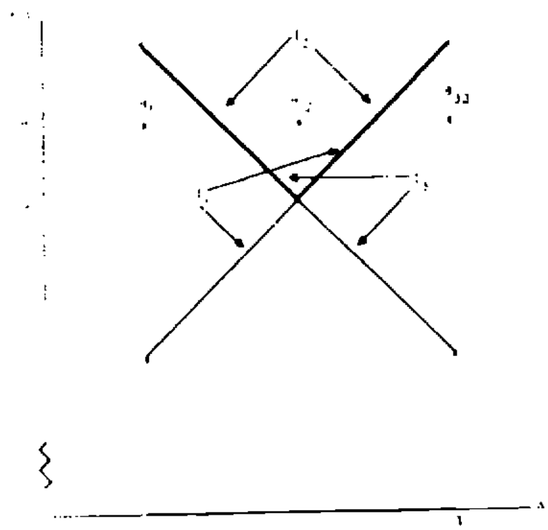

FIGURE 1

III. A PARTIAL ORDERING OF EQUILIBRIA

Before proceeding we need to introduce some additional notation. Defining an outcome as a mapping of individuals to action-price pairs, let Θ_U, Θ_I denote the sets of outcomes in ORUF and ORIF equilibria respectively.

Lemma 1: *In both the optimal reaction and the Nash equilibria when the informed move first, all contracts generate zero profits.*

This result is trivial. The uninformed observe the action of the informed and compete with one another driving their profits to zero.

Because of the possibility of nonzero profit contracts when the uninformed move first, it is not true for all economies that $\Theta_U \subset \Theta_I$ (see example 4 above.) However, if we impose the following restrictions on $c(a,i)$ and $\phi(a,i)$ we can eliminate ORUF equilibria in which some contracts generate non-zero profits.

A1. $\forall i \in I$, A_i is convex, $\phi(a,i)$ is continuous in A_i, and $c(\cdot,i)$ is continuously differentiable,

A2. For any pair of contracts $(\hat{a},\hat{p}),(\tilde{a},\bar{p})$ such that there is a type i that is indifferent between those contracts, then no type $j \neq i$ can also be indifferent between that same pair of contracts. (Note that this assumption is analogous to the usual assumption that indifference curves satisfy a single crossing property in a one dimensional action space.)

A3 If B denotes the union of the boundaries of A_i, then for $\forall a \in B$, $\forall i \in I$, $\phi(a,i) - c(a,i) < 0$.

Assumption A3 ensures that the equilibrium contract will lie in the interior of the action space of all agents. If a contract were on the boundary of the feasible set of actions of some agent then, notwithstanding A1 and A2, it might not be possible to attract that agent while repelling other agents who either are also choosing that contract or who are indifferent between that contract and the contract they are choosing in equilibrium.

Lemma 2. If A1, A2, A3 hold then each contract in an ORUF equilibrium generates zero profits.

Proof.

If $\forall a \epsilon \bar{A}$, $P(a) > \phi(a, f^{-1}(a,\pi))$, then firms generate losses and would be better off offering prices that did not elicit any trades. Thus $\exists a \epsilon \bar{A}$, such that $P(a) \leq \phi(a, f^{-1}(a,\pi))$. (Recall that $\bar{A}$ is the set of realized actions when the uninformed move first.)

Consider $a \epsilon \tilde{A}$ such that $P(a) < \phi(a, f^{-1}(a,\pi))$. Then, given the preferences of the informed and uninformed agents, that contract makes positive profits and hence is offered by every uninformed agent. For any $a \epsilon \tilde{A}$ such that $P(a) > \phi(a, f^{-1}(a,\pi))$ that contract is only offered by one uninformed agent: If two or more uninformed agents were offering the same money losing contract, one of those agents could increase its profits by lowering its price, so that no informed agents would purchase its contract, without affecting the distribution of actions chosen by the informed agents.

Thus, if some contracts make positive profits and others make losses, we would find all the positive profit contracts being offered by all the uninformed agents and each negative profit contract being offered by a single uninformed agent. The negative profit contract would offer the highest price consistent with preventing informed agents taking it from switching to another contract—the only reason negative profit contracts are offered is because of their sorting effects. Because uninformed agents can choose to offer prices which do not result in trades we know that in equilibrium all uninformed agents make non-negative profits. Assume some uninformed agents make positive profits equal to π^*. Consider an uninformed agent k whose profits are less than or equal to the average for the uninformed. That uninformed agent could perturb the equilibrium price schedule by offering a price schedule $\hat{P}(a) = P(a) + \epsilon(a)$, where $\epsilon(a)$ is everywhere positive but arbitrarily close to zero; the $\epsilon(a)$ function is chosen so that no informed agent chooses a different action from those induced by $P(a)$. (Our assumption that informed agents are risk neutral is sufficient to ensure the existence of such a function.) That contract enables the uninformed agent to capture the entire market and earn profits that are arbitrarily close to the average profits times the number of uninformed agents, which, of course, exceeds the average profits. Since we have chosen agent k such that its original profits were no greater than the average, this deviation is profitable for agent k. Therefore, an ORUF equilibrium cannot be characterized by any uninformed agent making positive profits.

Thus the only possibility left to consider is that each uninformed agent offers a combination of contracts some of which lose money while others make money. Each combination breaks even. The positive profit contracts are offered by all the uninformed agents. Each loss generating contract is offered by only one uninformed agent. Informed agents choosing money losing contracts are indifferent between that contract and at least one positive profit contract. Consider a contract $(\hat{p},\hat{a})$ that generates positive profits. (From A3, $\hat{a}$ lies on the interior of every set A_i.) Let S denote the set of types of agents who are indifferent between $(\hat{p},\hat{a})$ and the contract they are choosing, which could be $(\hat{p}, \hat{a})$. Let j denote the type of agent with the steepest indifference curve in the two dimensional surface $A^k \times P$ through $(\hat{p},\hat{a})$, where A^k is the k^{th} dimension of the action space. From A2 there is only one such type, and j's indifference surface through $(\hat{p},\hat{a})$ does not pass through a contract chosen by some type $i \epsilon J$, $i \neq j$.

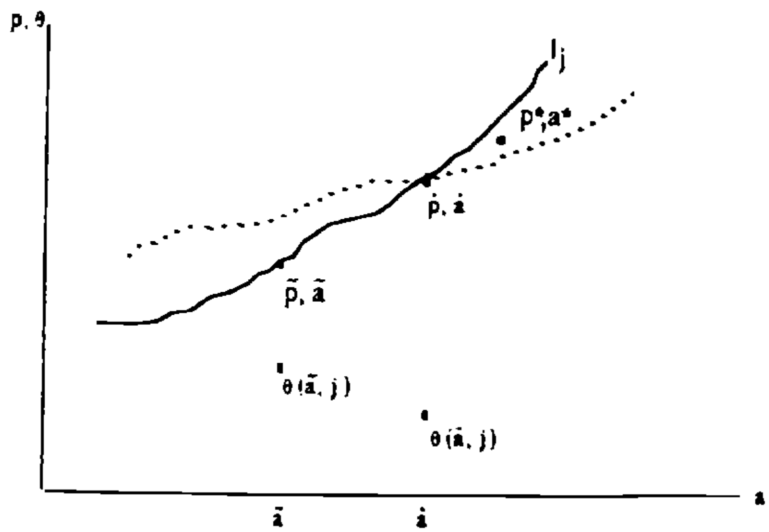

FIGURE 2

Suppose j is choosing $(\hat{p},\hat{a})$ and trade with that type is profitable. Then from A1, A2, A3, and the previous assumption that there is a finite number of types of agents, a contract can be offered in the neighborhood of $(\hat{p},\hat{a})$ that attracts only type j, and all of them. From A2 and the existence of several uninformed agent offering contract $(\hat{p},\hat{a})$, this new contract enables the uninformed agent offering it to make positive profits. Suppose j is choosing $(\hat{p},\hat{a})$ and trade with j is unprofitable for the uninformed. Then from A1, A2, A3, one of the uninformed could offer a price schedule with the following two properties: One of the contracts, (p^*,a^*), lies in the neighborhood of $(\hat{p},\hat{a})$ and is preferred to $(\hat{p},\hat{a})$ by every type in S except j and type j prefers $(\hat{p},\hat{a})$ to (p^*,a^*); and for every action $a \neq \hat{a}$ chosen by an informed agent, the new price schedule induces those agents to choose the same actions by offering a price that is slightly higher than that offered in the initial equilibrium. Thus the deviating firm attracts every type except j at an arbitrarily small change in prices and in the actions chosen by any agent. Since there is a discrete number of type j agents, each of whom generated losses, and since aggregate trade with the informed broke even, the new price schedule would make positive profits. It would approximate the profits and losses from all trades in the neighborhood of the old trades except those at (p^*,a^*) and would make discretely greater profits at (p^*,a^*) than at $(\hat{p},\hat{a})$. Finally, if j is choosing $(\hat{p},\hat{a})$ and trade with j breaks even, then consider the type of informed agents whose indifference curves through $(\hat{p},\hat{a})$ in $A^k \times P$ space has the second steepest slope, and proceed as before. If trade with that type also breaks even, continue until reaching a type with which trade generates non-zero profits.

Now suppose j chooses $(\bar{p},\bar{a}) \neq (\hat{p},\hat{a})$ and causes losses on contract $(\bar{p},\bar{a})$ for the uninformed agent offering that contract. There is a contract (p^*,a^*) in the neighborhood of $(\hat{p},\hat{a})$ that all informed agents in S except type j prefer to contract $(\hat{p},\hat{a})$ and type j prefers $(\hat{p},\hat{a})$ to this new contract. Any uninformed agent not offering $(\bar{p},\bar{a})$ could now offer this new contract and contracts in the neighborhood of all the old contracts except

$(\tilde{p},\tilde{a})$ so that all informed types continue to choose contracts in the neighborhood of their previous contracts. From A1 and the equilibrium condition that each uninformed agent breaks even, this new set of contracts generates positive profits. The new set of contracts approximates the sum of the profits and losses from all trades except those at contract $(\tilde{p},\tilde{a})$. Since trades at $(\tilde{p},\tilde{a})$ generated losses and the set of all previous trades generates zero profits, by omitting only trades with type j the new set of contracts generates positive profits.

Therefore, an equilibrium could not exist in which there are positive profit contracts because a new contract (or set of contracts) could be offered that generates positive profits for the agent offering that contract(s).

Theorem 1. Given A1, A2, A3, $\theta_U \subset \theta_I$.

Proof: From Lemmatta 1 and 2, regardless of the ordering of moves in an optimal reaction equilibrium, $\forall a \in \bar{A}$, the highest price being offered is equal to the expected value of a trade with a randomly selected agent choosing that action, and no informed agent will wish to deviate to a different contract within $\bar{A}$. All outcomes with these properties can be generated by Nash equilibria regardless of the order of moves in the game.[13] The maximum price offered for action $a \notin \bar{A}$ could be arbitrarily low in a Nash equilibrium of either game, thus sustaining any combination of moves with the two properties cited above. The outcome in θ_I has the additional restriction that deviations to $\hat{a} \notin \bar{A}$ are awarded contracts $p(\hat{a}) \geq \min_{i \in I} \phi(\hat{a},\{i\})$. Thus there is some bound on the penalty for departing the actions specified in ORIF that is not imposed on the penalties in a Nash equilibrium. This lower bound precludes some combination of plays which are supported in a Nash equilibrium from being supported as an ORIF equilibrium, but not conversely.

The penalties for deviations from actions specified in ORUF do not merely have a lower bound; they are precisely determined. If a price schedule $\hat{P}(a)$

13 There are also combinations of moves without these properties which are Nash equilibria—see examples 1 and 2.

$\neq$ P (a) is offered that induces some action $a \notin \tilde{A}$, that action is chosen by all agents (and only those agents) for whom

$$\hat{P}(a) - c(a,i) > P(f(i,\pi)) - c(f(i,\pi),i).$$

Since these reactions are known by all agents, in the ORUF equilibrium competition among the uninformed agents completely determines the contract offered for $a \notin \tilde{A}$: $\forall a \epsilon A$; $P(a) = \phi(a, f^{-1}(a,\pi))$.

Since $\phi(a,f^{-1}(a,\pi)) > \min_{i \epsilon I} \phi(a,i)$, threats in the ORUF to punish out-of-equilibrium behavior are less onerous than in an ORIF equilibrium. Since equilibrium outcomes are supported by these more onerous threats, $\theta_U \subset \theta_I$.

IV. REMARKS

1. Because the agents moving second could have strategies which depended on the observed distribution of moves of the agents moving first, Nash equilibria in these models may have some peculiar properties. In particular, if the uninformed move first, choosing price schedules before the informed choose actions, the strategies of the informed agents could enforce a positive profit Nash equilibrium.[14] The possibility of unreasonable positive profit Nash equilibria when the uninformed move first seems to be a basic feature of these models. If the informed move second, for them to choose the action which maximizes their expected payoff they must know the upper envelope of the price schedules of all agents. Hence the actions chosen by each informed agent is a function of this upper envelope. However, that condition would allow the informed agents to change their actions in such a way as to penalize an uninformed agent for *raising* his price(s), thus enforcing a positive profit Nash equilibrium when the uninformed move first. Obviously strategies of this form make little sense.

14 The converse is not true. When the informed move first, competition among the uninformed will result in their earning zero profits. This is because while peculiar strategies of the uninformed could force the informed to choose almost any actions, once those actions are chosen the uninformed will bid for the informed by offering prices that break even. There is no opportunity for either the informed or uninformed to use strategies that penalize high (or low) prices. We have assumed that after the informed choose their actions, they must trade with the firm(s) offering the highest price for their action.

2. The reader should also note that our definition of ORIF equilibrium is consistent with the Kreps-Wilson definition of sequential equilibrium in allowing extremely pessimistic beliefs about which agents took out-of-equilibrium actions. In Weiss(1983) more stringent restrictions were placed on the beliefs of the uninformed agents when faced with out of equilibrium actions. That paper defined a Robust Expectations Equilibrium (REE). It assumed that all agents with the same equilibrium strategy (and for whom the out-of-equilibrium action was feasible) are equally likely to have chosen a particular out-of-equilibrium action. Because the Weiss(1983) results are similar to ours, we suspect that $\theta_U \subset \theta_I$ holds for a broad class of definitions of equilibrium.

3. Many recent macro-economic models have been concerned with the effect of exogenous shocks on the equilibrium of an economy. If the exogenous shocks were observed by the informed but not by the uninformed, then, even if the uninformed moved first, optimal reactions would not be sufficient to enable the uninformed to know which informed agents would choose which contracts. The uninformed would not be able to offer contracts that separate the informed agents, since the preferences of the informed would be affected by these exogenous shocks that are unknown to the uninformed.

V. POSTSCRIPT

Since this paper was written in 1981 there has been considerable research done on game theoretic models of markets with asymmetric information. In this postscript we shall try to connect the approach we took with current research on models with asymmetric information and comment on the direction we think research in this field should be headed

Most research on the nature of equilibria in markets with informational asymmetries has focused on what we have called signaling games (Rosenthal and Weiss [1984], Wilson [1980] and Dasgupta and Maskin [1986] are exceptions). That research has been directed toward constructing new definitions of equilibrium that eliminate unreasonable Nash equilibria of the sort presented in the examples in section 2 of this paper, and more particularly, find a unique equilibrium. For instance Cho and Kreps show that by imposing the "intuitive criterion" on a particular formulation of the Spence signaling game there is a unique equilibrium outcome of the signaling game: the Pareto efficient separating equilibrium. Similarly Noldeke and van Damme obtain uniqueness in a different formulation of the Spence signaling game by imposing their "plausibility" criterion.

We think the quest for a single "correct" equilibrium notion that will predict the precise outcome for any signaling game is unlikely to be useful for understanding real economic problems. Details of the economic structure of the market being analyzed and the entire history of all relevant interactions are likely to have important effects on behavior in almost any interesting economic setting.

One reason for this is that, as we have seen, the nature of the equilibrium depends on inferences that participants make about out-of-equilibrium moves and the consequence actions which those moves give rise to. These inferences will be sensitive to the economic context in which the out-of-equilibrium move occurs, and the history of past interactions. The standard theory assumes that participants are rational, know that other participants are rational, know all the payoffs, and do not make mistakes. One inference that could be drawn from observing an action which a theory says should never occur is that some aspect of the theory is wrong. Which aspect of the theory is thought to be wrong will depend on the precise nature of the out-of-equilibrium move, and the economic context in which it took place. Since these inferences could favor the deviator, there will be situations in which it will be in the strategic interest of one participant to make what would appear, in the standard theory, to be an out-of-equilibrium move (though from a different perspective, a move which is fully rational.) For instance, it may be in the interest of one participant to be thought of as someone who frequently makes mistakes.

One way to address these issues is to formulate a model in which there are no out-of-equilibrium events. All feasible observations occur with some probability, either because of a rich heterogeneity in the types of participants, or because of a variety of kinds of errors.[15] With the latter approach the results may be highly dependent on the particular nature of errors introduced. Thus there may be no general theory (although there may be a general approach), since the sources and nature of errors in one market context may differ markedly from those in another. In general, participants will have a subjective probability distribution over the reasons for out-of-equilibrium moves.

There is a second reason that institutional and historical considerations will almost inevitably be drawn into a relevant analysis of equilibria. One of the lessons we have learned in the past decade is that the nature of the equilibrium of a game is highly dependent on the precise specification of the game. A slight reformulation of the standard Rothschild-Stiglitz insurance model can yield the quite different solution discussed by Wilson. Standard game theoretic models have participants forming beliefs and making inferences in an introspective manner; their reasoning is based on asking, what would a rational individual do in such a situation? But to make those inferences, two things are required: (a) each participant must know (or believe) that the other participants have precisely the same (correct) understanding of the rules of the game; and (b) there must be common knowledge of rationality, that is, each participant must not only believe that his opponent is rational, but that his opponent believes that he is rational, and that his opponent believes that he believes that his opponent is rational...Neither of these assumptions are reasonable, and it is certainly not reasonable to assume that all participants have confidence in the reasonableness of these assumptions.

15 Myerson (1979) focuses on mistakes in the evaluation of payoffs. Weiss (1983) and Simon (1987) focus on unintended actions.

Not being able to precisely model all the relevant details of any market interaction, and not trusting their judgments about the appropriateness of particular models, individuals tend to rely heavily on past experience to make inferences and judgments about their best course of action.

There are still other reasons that an historical-institutional approach is required. In many models, there are multiple equilibrium. There is no way, by introspection alone, that individuals can figure out what it is that their opponents are likely to do. History provides a natural coordinating mechanism for the choice of equilibria; unfortunately, history does not necessarily choose Pareto efficient equilibria.

Of course, even after we acknowledge the importance of history in selecting an outcome in the current period, questions remain as to how the past outcomes came about, under what conditions will a particular outcome persist, will future outcomes be near or far from those in the recent past, and how will the outcomes change over time. For instance will outcomes cycle? All of these questions are, of course, closely interconnected.

In the end, the strongest argument for an historical-institutional approach, is this: economics is a behavioral science; it is concerned with explaining a particular aspect of social behavior. In most contexts, individuals rely heavily on past experience to make inferences and judgments; they seldom rely exclusively, or even mainly, on introspective analysis. The question of why this is so remains a legitimate subject for enquiry.

It is perhaps worth noting that not only are our views supported by general observations of behavior, both of firms and individuals, but our views have also been widely confirmed by experimental evidence. Even in the simplest finitely repeated prisoner dilemma games, the predictions of standard game theory are not borne out. The persistence of—what appears from one perspective to be—unreasonable outcomes may be due to past learning or enculturation that teaches the players to value a particular process as opposed to payoff. For instance, if players are taught that cooperation is good, irrespective of the payoffs, then players might even cooperate in a one-shot prisoner's dilemma game.

Clearly whether or not a particular past interaction is relevant is a subjective decision of the participants in the market. It will be difficult for researchers to reach a consensus concerning which aspects of the past are relevant. However, because judging relevance is difficult does not imply that the entire history of past interactions is irrelevant. In particular, we would argue that the choices the uninformed made in response to actions chosen by a previous cohort of informed players will affect the actions the current cohort of informed players choose. (Note while we are very sympathetic to the forward induction arguments made by Kohlberg and Mertens on the importance of a particular player's past actions as signals of that player's future actions, we are making a different point.)

The points we have just made may be illustrated by the employment-education relationship which has motivated much of the literature on signaling models - including our own analysis. Currently, the standard treatment of that problem has become to consider a single informed agent whose type is randomly chosen from some distribution. The informed agent then chooses an education level (action) and the key question addressed by most research in this field is how firms react to out-of-equilibrium education levels (actions).

In practice, however, many individuals simultaneously choose whether or not to continue in school. In making these decisions they observe the wages offered to current and past school leavers at different education levels. Thus although not a repeated game, since both the players and the state variables change over time, informed and uninformed agents will infer from the outcomes of the previous period what payoffs they are likely to receive in the next period.[16]

16 Not only is there considerable experimental literature showing that equilibrium outcomes tend to persist, even if they are Pareto inefficient, but looking across countries it is hard to account for some of the variations in patterns of education in other than historical terms.

It does not strike us as a fruitful exercise to see what the outcomes of a particular game would be the first time it, or any similar, game were played. It seems to be a better operating assumption to think that history is always relevant and that there is never a first period. Consequently we would argue that ahistoric models (including our own) are seriously flawed, and that in trying to explain equilibrium outcomes researchers must not only take into account strategic considerations of the participants, but also the history of past play of this and similar games and any other relevant experiences of the participants. Indeed history (broadly defined to include learned notions of fairness) may even enforce what appear to the analyst to be unreasonable equilibria.

Similarly, the inferences drawn from out-of-equilibrium moves are likely to depend critically on the context. For instance if an equilibrium analysis suggests that no one should drop out of school within one week of graduation from high school nor should anyone pursue less than one year of junior college, the inferences that potential employers are likely to draw about a person who dropped out just prior to graduation are likely to be different from the inferences drawn about a person who went to college for one week. The latter is likely to have discovered he didn't like that college or college in general, the former is unlikely to have discovered one week prior to graduation that he so disliked high school that he didn't want to continue for the last week and graduate. In the case of the college dropout the unanticipated move can best be explained by misperceptions of one's own tastes. In the case of the high school dropout the unanticipated move may be best explained by irrational behavior or by some exogenous event.[17]

We have emphasized the difficulties of making judgments about the appropriateness of particular equilibrium concepts in the abstract; one must

17 In the high school attended one of the authors a student changed schools just prior to graduation because two of his fellow students, with whom he was not on friendly terms, were observed carrying guns.

analyze behavior within particular contexts. Individuals are seldom in the single-play, one period context envisaged in some standard formulations. This may explain why, when we place them in experimental situations corresponding precisely to those theoretical models, they so frequently behave in ways which are not consistent with the "theory." They extend to these highly stylized and unrealistic situations modes of behavior that were adapted to the more complex, dynamic environments in which they live.

Recent attempts to develop more dynamic models seem to us to represent one of the more fruitful lines of on-going research. While we have noted one aspect of this—the development of historic models, in which individuals use past experience to formulate their expectations—there are three others to which we would like to call attention.

First, in this paper, we have contrasted models in which the informed move first with those in which the uninformed move first. But the question of who moves first should not be exogenously imposed. We not only need to know which of these assumptions is more appropriate in various market contexts, but why.

Second, even in simple (one shot, no repeated play) markets, there may be complex dynamics. Elsewhere (Stiglitz-Weiss, 1987), we have provided one detailed example, in the context of the credit market, which we have analyzed as a four move game. Banks announce a set of policies; borrowers make applications; banks decide which of the applications to accept; and finally borrowers decide which of the loan offers they wish to accept. Equilibria in these multi-move (but single transaction period) games may be markedly different from those analyzed in the simpler games discussed in this paper. Deciding on the appropriate order of moves when transaction periods overlap is likely to affect the results of any analysis and is likely to require a detailed understanding of the market being analyzed.

Third, many of the actions involved in signalling and screening games take place over an extended period of time. They involve, as we have already noted, sequential decision making. And there may or may not be reversibility. Consider models in which firm's issuance of equity is used as a signal. If owners choose not to issue equity, it indicates that they believe that returns are high, and accordingly, potential buyers will pay more for their shares. But the original owners are, in general, not committed to retaining their ownership shares forever. Having sold some of their shares at a high price (because purchasers believed that they were going to retain their shares), they may subsequently sell more of their shares, at admittedly a lower price. Implicitly, earlier theories assumed that the original owners could make a commitment not to sell their shares in the future. In the absence of such a commitment, market equilibrium is markedly different from that characterized by the earlier models. (See Gale and Stiglitz (1986)).

Similarly, in the education market, individuals make decisions about whether to go to school for one more year on a year to year basis. The dynamic equilibrium may entail pooling, in contrast to the standard model, where individuals at birth commit themselves to a level of schooling.

The (partial) reversibility of some actions also introduces some inherent asymmetries into the choices of agents. An individual that drops out of school after 9th grade can later choose to resume her education. A college graduate cannot later choose to have had only 11 years of education. With imperfect capital markets consumption today precludes investing tomorrow, but saving today leaves open the option of saving or consuming tomorrow.

To sum up, we think that future research in the economics of information should have one or several of the following features:

1. It should be explicitly dynamic with stochastic state variables.
2. History should be allowed to affect the equilibrium outcomes.
3. Responses to out-of-equilibrium moves should depend on the institutional features of the market (including past interactions) and the nature of the particular out-of-equilibrium move.

BIBLIOGRAPHY

Banks, J. and J. Sobel. "Equilibrium Selection in Signaling Games." Econometrica 55 (1987): 647- 63.

Bhattacharya, S. "Nondissipative Signals and Dividend Policy." Quarterly Journal of Economics 95 (August 1980): 1-24.

Cho, I-K. and D. Kreps. "Signaling Games and Stable Equilibria." Quarterly Journal of Economics, (1987).

Dasgupta, P. and E. Maskin. "The Existence of Equilibrium in Discontinuous Economic Games, I:Theory." The Review of Economic Studies Vol. LIII, (1), No. 172, 1 (January 1986): 1-26.

________________. "The Existence of Equilibrium in Discontinuous Economic Games, II: Applications." The Review of Economic Studies Vol. LIII (1), No. 172, 1 (January, 1986): 27-42.

Fudenberg, D. and J. Tirole. "Sequential Bargaining with Incomplete Information." Review of Economic Studies 50 (April 1983): 221-248.

Gale, I. and J. Stiglitz. "Multiple Stock Offerings and the Financing of New Firms." Financial Research Center, Research Memo #73 (1986) Princeton University.

Guasch, J. and A. Weiss. "Wages as Sorting Mechanisms in Markets with Asymmetric Information." Review of Economic Studies 47 (July 1980): 653-64.

________________. "An Equilibrium Analysis of Wage-Productivity Gaps." Review of Economic Studies 49 (October 1982): 485-97.

Kohlberg, E. and J. Mertens. "On the Strategic Stability of Equilibria." Econometrica 50 (1986): 1003-1039.

Kreps, D. and R. Wilson. "Sequential Equilibrium." Econometrica 50 (July 1982): 863-94(b).

———— "Reputation and Imperfect Information." Journal of Economic Theory. 27 (August 1982): 253-79.

Kreps, D., Milgrom, P., Roberts, J. and R. Wilson. "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma." Journal of Economic Theory 27 (August 1982): 245-52.

Lazear, E. and S. Rosen. "Rank-Order Tournaments as Optimum Labor Contracts." Journal of Political Economy 89 (October 1981): 841-64.

Milgrom, P. and J. Roberts. "Limit Pricing and Entry Under Incomplete Information: An Economic Analysis." Econometrica 50 (March 1982): 443-459(a).

————. "Predation Reputation and Entry Deterrence." Journal of Economic Theory 27 (August 1982): 280-312(b).

Myerson, R. "Refinements on the Nash Equilibrium Concept." International Journal of Game Theory 7 (1978): 73-80.

Nalebuff, B. and J. Stiglitz. "Prizes and Incentives: Toward a General Theory of Compensation and Competition." Bell Journal of Economics 13 (Spring 1983): 131-47.

Noledeke, G. and E. von Damme. "Signalling in a Dynamic Labor Market." Bonn University Discussion Paper No. A148.

Riley, J. "Competitive Signaling." Journal of Economic Theory 10 (April 1975): 174-86.

————. "Informational Equilibria." Econometrica 47 (March 1979):331-360.

Rosenthal, R. and A. Weiss. "Mixed Strategy Equilibrium in a Market with Asymmetric Information." Review of Economic Studies 51 (April 1984): 333-42.

Rothschild, M. and J. Stiglitz. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Information." Quarterly Journal of Economics 90 (November 1976): 225-43.

Rubinstein, A. "Strong Perfect Equilibrium in Supergames." International Journal of Game Theory 9 (1979): 1-12.

Selten, R. "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games." International Journal of Game Theory 4 (1975): 25-55.

Spence, A. M. "Job Market Signaling." Quarterly Journal of Economics 87 (August 1973): 355-74.

———. "Market Signaling: Information Transfer in Hiring and Related Processes." Harvard University Press, Cambridge MA. (1974).

———. "Competition in Salaries and Prerequisites for Jobs." Quarterly Journal of Economics 90 (February 1976): 51-74.

Stiglitz, J. "Theory of Screening, Education and Distribution of Income." American Economic Review 65 (June 1975): 283-300.

———. "Incentives, Risk and Information: Notes Towards a Theory of Hierarchy." Bell Journal of Economics 6, #2 (Autumn 1975): 552-579.

———. "Information and Capital Markets." C. Cootner and W. Sharpe ed. Financial Economics: Essays in Honor of Paul Cootner. Prentice-Hall, Englewood, N.J., 1982.

Stiglitz, J. and A. Weiss. "Credit Rationing in Markets with Asymmetric Information." American Economic Review 71, #3 (June 1981):393-410.

Van Damme, E. Stability and Perfection of Nash Equilibria. Springer Verlag, Berlin (1987).

———. "Alternative Approached to Analyzing Markets with Asymmetric Information." American Economic Review 73 (March 1983: 246-49.

Weiss, A. "Education as a Test." Bell Laboratories Discussion Paper 101, (1977), forthcoming in Economics of Education Review.

———. "A Sorting-cum-Learning Model of Education." Journal of Political Economy 91 (June 1983): 420–42.

Wilson, C. "A Model of Insurance Markets with Incomplete Information." Journal of Economic Theory 16 (1978): 167-207.

———. "The Nature of Equilibrium in Markets with Adverse Selection." The Bell Journal of Economics 11, #1 (Spring 1980a): 108-130.

Yabushita, S. "Theory of Screening and the Behavior of the Firm: A Comment." American Economic Review 73 (March 1983): 242–45.