Institute for Mathematics and Its Applications

BAYESIAN ANALYSIS IN ECONOMICS AND GAME THEORY

WORKSHOP ON BAYESIAN ANALYSIS IN GAME THEORY AND ECONOMICS

The abstracts in this collection demonstrate some of the depth and range of modern research in Bayesian models. The field covered has grown so rapidly that it is no longer possible to cover all of it with any thoroughness in a single conference. The most notable omissions are abstracts in the areas of search and learning, which have played an increasingly significant role in explaining economic phenomena. Offsetting these omissions is an unusually thorough coverage of the implications of Bayesian incentive theory and the inclusion of exciting new analyses of the role of bounded rationality in economic organization theory.

The collection opens with two abstracts by Myerson introducing Harsanyi's Bayesian game model and the closely related issues of Bayesian incentive compatibility, with an application to bargaining problems. Kreps and Harsanyi then contribute their ideas about refinements of the underlying noncooperative equilibrium concept. Kreps's paper focuses on "stable" equilibria, and characterizes their nature and efficiency properties in a class of signalling games based on Spence's labor market model. Harsanyi applies the Harsanyi-Selten equilibrium point selection criterion to Bayesian games.

The next two papers challenge the applications and conclusions derived from Bayesian game models. Crawford argues that studies of Bayesian bargaining games finesse issues of the communication that would be required to reach any agreement. Ledyard adds that the auxillary assumptions used by Bayesian modelers are at least as important as the incentive issues themselves in leading to the predictions attributed to the theory.

Rob and Samuelson apply the Bayesian incentive approach to the problem of bargaining costs. They find that some costs are inevitable, and that these substantially undermine the usefulness of the "Coase Theorem". How can one reconcile the notion that efficient bargains in the traditional sense are unachievable with the idea that market allocations are efficient in an Arrow-Debreu world? Gresik and Satterthwaite show by an example that inefficiencies (or "bargaining costs") as a fraction of the potential gains from trade may dwindle rapidly as the number of traders grows larger.

The next three abstracts turn from the theme of incentive compatibility to examine the implications of imperfect information in financial economics. Glosten and Milgrom study how the bid-ask spread in a security market can be generated by the need of the market maker to earn back money he loses on trades with insiders. Harris and Raviv study the decision of a firm to convert its callable debt, given that its managers have better information than the market as a whole. Dubey, Geanakopolos, and Shubik develop a tractable game model of trading in general equilibrium with private information that offers an alternative to the rational expectations equilibrium theory that makes the process of price formation explicit.

There follow two applications of Bayesian games to the field of Industrial Economics. Fudenberg and Tirole study the decision by a firm to exit from a product market. The exit decision is based on the inference that is drawn over time about the nature of competition in the market. Milgrom and Roberts consider the problem of introducing a new product of high quality to skeptical consumers. According to their analysis, high introductory prices, low introductory prices, and advertising can each be part of a producer's strategy to signal its high quality, depending on its production costs.

The final two abstracts point to a possible new direction of research in Bayesian economic models. These concern the use of information in organizations when individuals make mistakes or are limited in their information processing capacity. Sah and Stiglitz consider how the structure of an organization compounds or corrects the errors made by decision makers at each level. They find that in certain environments that they deem to be realistic, "polyarchies", in which decisions are substantially decentralized, perform better than "hierarchies", in which all decisions must be approved by the central authority. Geanakopolos and Milgrom explain the decentralization of decisions as a response to bounded rationality and they characterize hierarchies as mechanisms for coordinating decentralized decisions. In their analysis, the delays imposed by decision making in a hierarchy limit its depth and scope.

Paul Milgrom
Roger Myerson
Organizers

BAYESIAN EQUILIBRIUM AND INCENTIVE-COMPATIBILITY: AN INTRODUCTION

by

Roger B. Myerson

J. L. Kellogg Graduate School of Management
Northwestern University
Evanston, Illinois 60201

Abstract

This paper is an introduction to the analysis of Bayesian games with incomplete information. First, the logical foundations of the Bayesian model are discussed. Following Harsanyi [1967-8] and Mertens and Zamir [1982], it is argued that the Bayesian games are appropriate models for general situations of conflict and cooperation under uncertainty. To describe rational behavior in Bayesian games, two solution concepts are presented: Bayesian equilibrium, for games in which the players cannot communicate; and Bayesian incentive-compatibility, for games in which the players can communicate. The set of incentive-compatible mechanisms is a convex set and is often easier to analyze than the set of equilibria without communication.

In the second half of the paper, techniques are developed to characterize the efficient incentive-compatible mechanisms. The most important new concept in this analysis is virtual utility, which differs from real utility by taking into account the costs of satisfying incentive constraints. With this conceptual structure, we can make general statements about the nature of optimal social systems when individuals have difficulty trusting each other. The cases of pure moral hazard problems and pure adverse selection problems are considered in some detail.

References

J. Harsanyi [1967-8], "Games with Incomplete Information Played by 'Bayesian' Players," Management Science 14, 159-189, 320-334, 486-502.

B. Holmstrom and R. Myerson [1983], "Efficient and Durable Decision Rules with Incomplete Information," Econometrica 51, 1799-1819.

J. F. Mertens and S. Zamir [1982], "Formaliziation of Harsanyi's Notions of 'Type' and 'Consistency' in Games with Incomplete Information," C.O.R.E. discussion paper, Université de Louvaine.

R. Myerson [1979], "Incentive Compatibility and the Bargaining Problem," Econometrica 47, 61-73.

R. Myerson [1983], "Mechanism Design by an Informed Principal," Econometrica 51, 1767-1797.

R. Myerson [1984], "Two-Person Bargaining Problems with Incomplete Information," Econometrica 52, to appear.

Signalling Games and Stable Equilibria

David M. Kreps
Stanford University

In many games with incomplete information, one party sends a signal about privately held information to a second party, who thereupon takes an action that affects both. Examples include Spence's model of job market signalling, signalling about the quality of a product, signalling about one's intentions or one's cost structure, and sequential bargaining. Because off-the-equilibrium-path beliefs are unrestricted in sequential equilibria, such games typically possess many sequential equilibria. The number of equilibria can, however, be greatly reduced through restrictions on off-the-equilibrium-path beliefs. A number of authors (Grossman, Kreps and Wilson, Milgrom and Roberts, Rubinstein) have used a particular intuitive argument of the form: One type of the first party should be able to distinguish himself from other types if that type can send a message that other types would "never profitably" send.

In this paper, I illustrate the power of this intuitive argument by showing that, in Spence's model of job market signalling, it gives a unique equilibrium outcome. Then I connect this argument with Kohlberg-Mertens' stability: For generic extensive signalling games, satisfaction of this intuitive argument is necessary for an outcome to be a stable equilibrium outcome.

"A Bayesian Criterion for Equilibrium-Point Selection"

ABSTRACT

by John C. Harsanyi
University of California, Berkeley

Use of Bayesian ideas in modelling games with incomplete information is well known. [Harsanyi, 1967-68] Here, I would like to discuss another use of Bayesian ideas in game theory.

In recent years many researchers have more or less independently reached the conclusion that, for many purposes, game situations can be best modelled as noncooperative games--if we characterize a noncooperative game as a game whose solution is defined in terms of its equilibrium points [Nash, 1951] or, preferably, in terms of its perfect equilibrium points [Selten, 1975].

Yet, noncooperative-game models immediately pose the problem that almost all nontrivial games have very many, often infinitely many, essentially different perfect equilibrium points. This may be called the multiplicity problem.

Rubinstein [1982] has shown that in many cases this problem can be avoided by assuming that the players apply time discounts to payoffs they expect to receive only at later stages of the game. This is certainly a very ingenious approach but seems to be rather ad hoc. In particular, it is doubtful that the players will significantly discount a payoff they expect to obtain after a relatively short period of bargaining, which may last only a few days or even a few minutes.

A different approach to the multiplicity problem has been developed by Reinhard Selten and me. It involves providing general mathematical criteria--ones based largely on Bayesian decision theory--for selecting one specific equilibrium point of any given game as the solution. [Our theory will be described in a forthcoming book by the two of us, A General Theory of Equilibrium Selection in Games. For a brief preliminary outline of our theory, see Harsanyi, 1982.]

Suppose the players want to decide which one of two otherwise admissible equilibrium points $s^* = (s_1^*, \ldots, s_n^*)$ and $s^{**} = (s_1^{**}, \ldots, s_n^{**})$ should be the solution. Then, under our theory, they will choose s^* as the solution if

(1) s^* payoff-dominates s^{**}, in the sense that s^* yields a higher payoff $H_i(s^*) > H_i(s^{**})$ to every player i. (Thus, payoff dominance is the same as strong Pareto superiority). Or if

(2) neither equilibrium point payoff-dominates the other, yet s^* risk-dominates s^{**}.

For lack of space I cannot define the concept of risk dominance here and have to refer the reader to op. cit., pp. 183-187. But let me say that risk dominance is meant to be a formalization of the Bayesian subjective-probability and the Bayesian expected-payoff considerations that all of us use informally in everyday life when we choose between two alternative equilibrium strategies.

For example, suppose that A wants to sell his car to B and wonders whether he should ask $3,000 or $4,000 for it. In the end, he decides to ask only $3,000. Why will he do this? It is clear on a common sense level that the reason must be his feeling that the probability of B's paying $4,000 for the car is so low that his expected utility payoff from the strategy of asking only $3,000 is higher than that from the strategy of asking $4,000.

Or, suppose that A is negotiating, not with one potential buyer, but rather with k such individuals. Common sense tells us that, other things being equal, A can now reasonably try for a higher price because, even if the probability ρ that any given buyer should pay this price is rather low, the probability

$$\rho^* = 1 - (1 - \rho)^k$$

that at least one of them will do so may be quite high (as long as k is large enough).

Through our concept of risk-dominance, such Bayesian subjective-probability considerations play an essential role in our theory [cf. op. cit., pp. 261-265], as they do in any common sense analysis of strategic decision problems. In our view, it was a serious shortcoming of classical game theory that it excluded all such considerations from its conceptual framework.

To conclude, let me briefly describe the solution our theory selects for a two-person bargaining game with incomplete information on both sides. [For details, see Harsanyi, op. cit., pp. 266-279.]

Players I and II have to divide $100. Both have linear utility functions for money. If they cannot agree then they will obtain only their conflict payoffs c_I and c_{II}. Either conflict payoff can have the value 0 or the value α, where $0 < \alpha < 50$. Each of the four possible conflict-payoff vectors $(0, 0)$, $(0, \alpha)$, $(\alpha, 0)$, and (α, α) will occur with probability 1/4. Both players know this probability distribution and know their own conflict payoff. But neither knows the other player's conflict payoff. Each player i is called weak if $c_i = 0$ and is called strong if $c_i = \alpha$.

We assume that the game can have at most two stages, in which the players may state their payoff demands. But it may end--with an agreement or with a conflict--already at stage 1.

Our theory selects the following solution for this game: If $\alpha \leq 33\ 1/3$ then each player, whether he is strong or weak, will accept the payoff $u_i = 50$. Accordingly, the players will agree on these payoffs $u_1 = u_2 = 50$ already at stage 1.

In contrast, if $\alpha > 33\ 1/3$ then any strong player will insist on a payoff demand of $u_i = 50 + (\alpha/2)$. As a result, if both players are strong, they will cause a conflict already at stage 1. But if both players are weak, or if one is strong and one is weak, then they will always reach an agreement at stage 2. More particularly, since a strong and a weak player will act differently at stage 1, by stage 2 both players will know whether the other player is strong or weak. As a result, if both players are weak, they will agree on the payoffs $u_1 = u_2 = 50$; whereas, if one is strong while the other is weak then the former will obtain $u_i = 50 + (\alpha/2)$ while the latter will obtain only $u_j = 50 - (\alpha/2)$.

Note that these are exactly the payoffs they would obtain under the Nash solution if the game were one with complete information, in which both players knew each other's conflict payoffs from the start. Since at stage 2 our game does in fact involve complete information (as a result of information obtained at stage 1), I feel this is a very desirable property of our solution--a property that most other solution concepts for bargaining games with incomplete information fail to possess.

References

J.C. Harsanyi, "Games with Incomplete Information Played by Bayesian Players." Management Science, 14 (1967-68), pp. 159-182, 320-334, and 486-502.

"Solutions for Some Bargaining Games under the Harsanyi-Selten Solution Theory." Mathematical Social Sciences, 3 (1982), pp. 179-191 and 259-279.

J.F. Nash, "Noncooperative Games." Annals of Mathematics, 54 (1951), pp. 286-295.

A. Rubinstein, "Perfect Equilibrium in a Bargaining Model." Econometrica, 50 (1982), pp. 97-109.

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

JCH/rlw

Abstract of "Efficient and Durable Decision Rules: a Reformulation"
by Vincent P. Crawford, University of California, San Diego

This paper studies the limits of contracting as a method for achieving efficient resource allocation, with particular attention to the effects of informational asymmetries. When information is perfect, this question can be given a simple answer. Assume that agents must play a game, but that before play they can make complete, perfectly enforceable agreements about how to play the game. When such contracts are costless to make, arguments in the spirit of the Coase "Theorem" [1] suggest that rational agents can agree on a joint strategy choice that yields an individually rational, Pareto-efficient allocation. When agents have symmetric information and complete contingent contracts are possible, the situation is analogous to the perfect-information case and the perfect-information argument yields the same conclusion, with individual rationality and Pareto-efficiency properly reinterpreted, in their ex ante senses.

Significant new difficulties emerge when information is asymmetric. Then, it seems natural to evaluate the limits of contracting by assuming that agents can make complete, perfectly enforceable agreements about the mechanism that will be used to control their incentives when they play the underlying game, and that they evaluate those agreements according to their rational expectations of the effects of the incentives they create. Thus, the choice of a mechanism is analogous to the choice of an allocation under perfect information. Two cases can be distinguished. In the first, agents can make binding agreements before observing their private information. This case is essentially equivalent to the perfect- and symmetric-information cases once feasibility and efficiency are redefined, as in Myerson [3], to reflect incentive constraints. It can therefore be argued that in this case agents will agree on an ex ante incentive-efficient mechanism; this is the kind of generalization of the perfect- and symmetric-information results one would expect to obtain.

The second case, where agents cannot make binding agreements until after they have observed their private information, is considered in a recent paper by Holmstrom and Myerson [2]. This case is difficult to analyze because the process of mechanism design must, to yield good results, aggregate information about agents' preferences over mechanisms. Since these preferences are generally influenced by agents' private information, mechanism design at this stage may "leak" agents' private information, altering the incentives created by the chosen mechanism. Even if the effects of this information leakage are rationally anticipated, they can make it impossible for mechanism design at this stage to yield an incentive-efficient result. Holmstrom and Myerson provide an example to illustrate this problem, and propose a unanimous-voting characterization of the set of outcomes that agents can achieve despite the problems caused by information leakage in the mechanism-design process; they call such outcomes durable decision rules.

The present paper proposes a reformulation of Holmstrom and Myerson's concept of durability, in the spirit of noncooperative bargaining theory. Endogenous mechanism design is viewed as taking place within a two-stage game. In the first stage, agents bargain (with exogenous bargaining rules) over the mechanism that will be used to govern play in the underlying game; in the second stage they play the underlying game, responding to the incentives created by the agreed-upon mechanism, taking into account any information leaked in the process of agreeing on it. Perfect equilibrium outcomes of this two-stage process correspond to (but are generally different from) Holmstrom and Myerson's durable decision rules. It is argued that the rules which govern bargaining over mechanisms must be exogenous, at some level, for a game-theoretic analysis of endogenous mechanism design to be possible. Two sample specifications of these rules are proposed and discussed. Both use only common-knowledge information, are frictionless, and yield the "correct" answers when information is perfect or symmetric, or when information is asymmetric but mechanism design can take place before agents observe their private information. But one specification yields nontrivial durability restrictions when mechanism design takes place after private information is observed, while the other eliminates all such restrictions, making it possible for agents to agree on the most efficient mechanisms that are attainable when mechanism design takes place before agents observe their private information. The latter specification allows agents to create their own veil of ignorance, by endogenously restricting the language they can use to bargain over mechanisms, leaving it rich enough to allow the selection of efficient mechanisms, but coarse enough so that agents cannot benefit from proposals that leak their private information. The paper concludes with a discussion of the implications of these results.

References

[1] R.H. Coase, "The Problem of Social Cost," Journal of Law and Economics 3 (1960), 1-44.

[2] B. Holmstrom and R. Myerson, "Efficient and Durable Decision Rules with Incomplete Information," Econometrica 51 (1983), 1799-1819.

[3] R. Myerson, "Incentive Compatibility and the Bargaining Problem," Econometrica 47 (1979), 61-73.

THE SCOPE OF THE HYPOTHESIS OF BAYESIAN EQUILIBRIUM

John Ledyard

California Institute of Technology

There has recently been successful use of the hypothesis that a coordination mechanism be incentive compatible, in deriving restrictions on the form of "optimal mechanisms", by Myerson and Satterthwaite, Gresik and Satterthwaite, and Wilson. Similarly, diverse economic behaviors have been rationalized as equilibria of Bayesian games. The unanwered question by this research is whether it is the assumption of Bayes equilibrium behavior or the assumptions of the specific utility functions and prior beliefs which drive the results. In this paper, I try to identify those coordination mechanisms and performance functions for which there exists some set of utilities and beliefs such that that performance is the result of Bayesian Equilibrium behavior in that mechanism and in that environment. That is, I am interested in identifying which observed behavior can be rationalized as the Bayesian Equilibrium of some game. Economists will recognize a close similarity in spirit to the question, "what can be an aggregate excess demand function?"

We first apply a variation of the revelation principle to show that, subject to slight informational restrictions, there is an environment which rationalizes the given performance in a specific game form if and only if the performance function is itself an incentive compatible direct revelation mechanism. Thus, we need only answer the question, for what direct revelation mechanism is there an environment in which that mechanism is incentive compatible? Knowing that, we will know what observed behavior is consistent with the hypothesis of Bayes equilibrium.

Without further restrictions, any performance function, a function from types to outcomes, can be rationalized in a non-trivial way as a game equilbrium. Even if both preference orders, on outcomes, and prior beliefs are specified for each vector of types and each agent, almost any performance can still be rationalized by the selection of utilities consistent with the specified preferences. Only performance which is perfectly perverse for those preferences, when truth is always the least preferred strategy, can not be rationalizd.

It is only restrictions on utilities which limit the types of performance which can be consistent with Bayesian equilibrium behavior. Some indication of these restrictions is provided, but a full characterization remains to be done. It is shown that this problem is closely related to that of Savage _The Foundations of Statistics_ to describe beliefs and utilites such that specified behavior is rationalized by expected utility maximization.

In summary, even if the preferences and prior beliefs of all agents, as well as the game form, are pre-specified _any non-perverse performance can be rationalized as the result of Bayesian equilibrium behavior_. Only when severe retrictions are placed on utility functions can certain performance be eliminated as not being sensible.

The Coase Theorem - An Informational Perspective

Rafael Rob

University of Minnesota

March 1984

Abstract

A reexamination of the so-called Coase theorem when transaction costs are recognized is presented. When the latter are modeled as private information about individual damages we show that Coase's efficiency claim losses much of its force. It is perfectly possible that negotiations over suitable compensations might lead to a misallocation of resources. Under the assumption that the number of individuals is large these inefficiencies increase in magnitude. It appears that a substantial portion of the surplus is dissipated by bargainings.

0. Summary

The purpose of this paper is to rigorously examine the notions advanced by R. Coase in his fundamental paper "The Problem of Social Cost". The analysis is carried out in an adverse selection theoretic framework where transation costs are formalized as private information about individual damages. Given this interpretation of transaction costs and using the methodology introduced by Myerson [13], I will show that Coase's efficiency claim losses much of its force, i.e. that reaching a suboptimal outcome is a realistic possibility. More importantly, when the decision involves a large number of participants, these informational inefficiencies become rather severe which in turn, renders the scheme practically useless. Furthermore, the assignment of property rights has no bearing upon this market failure. Notwithstanding this, my study does not ratify the imposition of Pigourian taxes, or a centralized system of pollution control or any other regulatory measures. This is so since a governmental agency is subject to at least as many incentive contraints as the private sector, so that there is no a priori reason to believe that a central authority can improve upon an informational market failure. This analysis does invite, however, a further invetigation of the actual mechanisms customarily employed to resolve externality-type problems. The tradeoff which seems to emerge is between high average performance (with arbitration or voting schemes) and the preservation of the status quo ante for all individuals

The section below surveys the issues analyzed in this paper. In section 2 a formulation of the problem follows. Section 3 studies the model, derives closed-form solutions and provides and example illustrating the results. Section 4 is a quantitative examination of the welfare performance of the mechanism when the number of residents is large. Section 5 concludes the paper.

1. Introduction

A cornerstone of welfare economics is that externalities generate market inefficiencies. This has been variously attributed to fundamental non-convexities (Starrett [22]) or to the lack of markets for commodities which are not property in law (Arrow [1]). The basic problem with external effects is that they introudce an inherent divergence between private and social costs. Factories overpollute the environment simply because they are not held liable for the external costs they impose on others. In the absence of hazaradous waste regulation individual residents are forced to bear some of the costs of production which, in turn, makes firms' operations more profitable. Production and pollution levels consequently increase beyond what would appear to be socially optimal.

Having diagnosed the basic problem, a natural policy recommendation emerged (Pigou, [16]). Effluent-producing firms can be induced to act in the best

A Comment on the Coase Theorem

William Samuelson
Boston University

Abstract

Beginning with the observation of Ronald Coase, it has long been held that private bargaining can provide an antidote to the inefficiencies caused by externalities. Coase [1960] argued that

i) a pair of agents, by striking a mutually advantageous agreement, would obtain an efficient economic solution to the externality, and

ii) a change in the assignment of property rights or in the liability rule would not affect the attainment of efficient agreements.

The Coase "theorem" relies on a number of assumptions, some explicit, some implicit, among these that: agents have perfect knowledge of the economic setting including each other's utility function; in the absence of transaction costs, the agents will strike mutually beneficial agreements; and there exists a costless mechanism (a court system) for enforcing such agreements.

The presumption that the bargainers have perfect knowledge of and pursue mutually beneficial agreements -- assumptions borrowed from the theory of cooperative games -- is crucial for the Coase results. The usual argument is that rational bargainers would (should) never settle on a given set of agreement terms if instead they could agree on alternative terms which were preferred by both sides. The conclusion, according to this argument, is that any final agreement must be Pareto optimal.

While this reasoning seems compelling, it leaves a number of questions unanswered. By what bargaining procedure do the individuals actually arrive at a Pareto efficient agreement? Will alternative procedures attain such agreements? The usual presentation of the Coase theorem omits the specifics of the bargaining process. The presumption is simply that such agreements can and will be reached since it is in the joint interest of the parties to do so.

Though this conclusion is appealing, it is considerably stronger than the economist's customary hypothesis of individual rationality. In particular, it is unrealistic to suppose that the bargaining setting is one of perfect information. To a greater or lesser degree, each party will be uncertain about key aspects of the bargaining situation -- possibly about its own payoffs for alternative agreements and almost certainly about the other side's potential payoffs. In the case of an externality, two firms (one a polluter and one a victim of the pollution) might be engaged in negotiations aimed at a mutually beneficial agreement. Each firm would be knowledgable about its own cost (clean-up cost or pollution cost) but would have limited information about the costs of the other side.

The presence of limited or imperfect information is an impediment to the attainment of efficient agreements. Before the fact, neither side will know whether and what mutually beneficial agreements are available. Such agreements when they exist, must be uncovered by the bargaining process itself. Thus, the key question is whether there exists a suitable bargaining method, conducted by self-interested individuals, which can always achieve Pareto-efficient agreements.

This paper addresses this question by exploring the Coase Theorem in a setting of incomplete information. In contrast to the traditional version of the Coase Theorem, the results under incomplete information are as follows:

1) The parties affected by an externality will, in general, be unable to negotiate efficient aggrements all of the time.

2) The degree of efficiency of a negotiated agreement depends on

   a) which party is assigned the specific property right and

   b) the bargaining process employed by the parties.

3) Efficiency can be increased by allocating the property right via competitive bid, rather than relying on some pre-assignment.

The first proposition indicates that the presence of incomplete information imposes a second-best negotiation solution -- one in which ex post efficiency cannot always be attained. Once an intitial rights assignment is made, bargaining cannot insure the transfer of the right to the party that values it most highly. In this sense, the presence of incomplete information creates a kind of trading friction or a barrier to trade. The second proposition follows as an immediate consequence. Given the informational barriers to trade, the ultimate solution to the externality depends directly on the initial assignment of the right and on the bargaining method adopted by the parties. The last proposition indicates a potential remedy: efficiency can be increased by foregoing the rights assignment in the first place. Instead, the right is allocated via competitive bid (so that the agent willing to pay the most for the right obtains it) without recontracting.

The main propositions in this paper are applications or extensions of recent results in the area of resource allocation under uncertainty. Despite their importance, many of these results are not well-known (or if known, not assimilated into mainstream economic thinking). Thus, the conventional wisdom holds that in the absense of transaction costs (caused, for instance by numerous interested parties), bargaining can solve in principle the externality problem. But whether bargaining is an appropriate remedy depends on the kind of market failure present. If external effects are the sole cause of the failure, a bargaining solution which, in effect, internalizes the externality will be appropriate. But in the common case when external affects are accompanied by the presence of imperfect information, a property rights assignment followed by private bargaining will not be a fully-efficient remedy.

The first proposition indicates that the presence of incomplete information imposes a second-best negotiation solution -- one in which *ex post* efficiency cannot always be attained. Once an intitial rights assignment is made, bargaining cannot insure the transfer of the right to the party that values it most highly. In this sense, the presence of incomplete information creates a kind of trading friction or a barrier to trade. The second proposition follows as an immediate consequence. Given the informational barriers to trade, the ultimate solution to the externality depends directly on the initial assignment of the right and on the bargaining method adopted by the parties. The last proposition indicates a potential remedy: efficiency can be increased by foregoing the rights assignment in the first place. Instead, the right is allocated via competitive bid (so that the agent willing to pay the most for the right obtains it) without recontracting.

The main propositions in this paper are applications or extensions of recent results in the area of resource allocation under uncertainty. Despite their importance, many of these results are not well-known (or if known, not assimilated into mainstream economic thinking). Thus, the conventional wisdom holds that in the absense of transaction costs (caused, for instance by numerous interested parties), bargaining can solve in principle the externality problem. But whether bargaining is an appropriate remedy depends on the kind of market failure present. If external effects are the sole cause of the failure, a bargaining solution which, in effect, internalizes the externality will be appropriate. But in the common case when external affects are accompanied by the presence of imperfect information, a property rights assignment followed by private bargaining will not be a fully-efficient remedy.

The Number of Traders Required to Make a Market Competitive: The Beginnings of a Theory

Thomas A. Gresik and Mark A. Satterthwaite
Northwestern University

Model. The market consists of N identical objects, N sellers who each own one of the objects, M buyers who each seek to buy one of the objects, and money.

x_i Buyer i's reservation valuation of the object. It is known to i and unobservable to all others. Sellers and the other buyers regard x_i as distributed with positive density $f_i(\cdot)$ over the interval $[a_i,b_i]$.

z_j Seller j's reservation value. Buyers and other sellers regard it as distributed with positive density $h_j(\cdot)$ over $[c_j,d_j]$.

$F_i(\cdot)$ and $H_j(\cdot)$ are the distribution functions of the densities f_i and h_j respectively. The pair (F,H), where $F = (F_1, \ldots, F_M)$ and $H = (H_1, \ldots, H_N)$, is called the trading problem.

$g(x,z) = \Pi_{i=1}^{M} f_i(x_i) \cdot \Pi_{j=1}^{N} h_j(z_j)$ is the joint density of the reservation values where $x = (x_1, \ldots, x_M)$ and $z = (z_1, \ldots, z_N)$.

$g(x_{-i},z) = g(x,z)/f_i(x_i)$ and $g(x,z_{-j}) = g(x,z)/h_j(z_j)$ respectively describe the distributions of reservation values buyer i and seller j perceive themselves as facing where $x_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_M)$ and $z_{-j} = (z_1, \ldots, z_{j-1}, z_{j+1}, \ldots, z_N)$.

$p_i(x,z)$ and $q_j(x,z)$ are respectively the probabilities of an object being assigned to buyer i and seller j in the final distribution of goods.

$r_i(x,z)$ and $s_j(x,z)$ are respectively the payments to buyer i and seller j, e.g., $r_i < 0$ indicates that buyer i pays negative r_i units of money for receiving one unit of the object with probability p_i.

(p, q, r, s) is the 2M+2N vector of probability and payment schedules that constitute a trading mechanism.

Common knowledge is assumed concerning the joint density g and the trading mechanism's rules (p, q, r, s).

Balance of goods and balance of receipts constraints are: for all x and z

$$\sum_{i=1}^{M} p_i(x,z) + \sum_{j=1}^{N} q_j(x,z) = N, \tag{1}$$

$$\sum_{i=1}^{M} r_i(x,z) + \sum_{j=1}^{N} s_j(x,z) = 0. \tag{2}$$

Construction of an Ex Ante Efficient, Individually Rational, Incentive Compatible Trading Mechanism.

Virtual reservation values of buyer i and seller j are

$$\psi_i^B(x_i,\alpha) = x_i + \alpha \cdot \left(\frac{F_i(x_i) - 1}{f_i(x_i)}\right), \quad i=1,\ldots,M, \tag{3}$$

and

$$\psi_j^S(z_j,\alpha) = z_j + \alpha \cdot \frac{H_j(z_j)}{h_j(z_j)}, \quad j=1,\ldots,N, \tag{4}$$

where α is a nonnegative scalar.

$\psi(x, z, \alpha) = \{\psi_1^B(x_1\ \alpha), \ldots, \psi_N^S(z_N, \alpha)\}$; is the M+N vector of the traders' virtual reservation values.

$R_i(x, z, \alpha)$ and $R_j(x, z, \alpha)$ are the ranks of the elements $\psi_i^B(x_i, \alpha)$ and $\psi_j^S(z_j, \alpha)$ within ψ.

Define an α-mechanism as follows. Let:

$$p_i^\alpha(x,z) = \begin{cases} 1 & \text{if } R_i(x, z, \alpha) \leq N \\ 0 & \text{if } R_i(x, z, \alpha) > N \end{cases} \quad i=1,\ldots,M; \tag{5}$$

$$q_j^\alpha(x,z) = \begin{cases} 1 & \text{if } R_j(x, z, \alpha) \leq N \\ 0 & \text{if } R_j(x, z, \alpha) > N \end{cases} \quad i=1,\ldots,N. \tag{6}$$

Let $p^\alpha = (p_1^\alpha, \ldots, p_M^\alpha)$ and $q^\alpha = (q_1^\alpha, \ldots, q_N^\alpha)$. Any mechanism that has probability schedules p^α and q^α is called an α-mechanism.

The constraint on (p, q) that guarantees the mechanism's incentive compatibility and individual rationality is:

$$G(\alpha) = \int\ldots\int\{\sum_{i=1}^{M}\psi_i^B(x_i,1)p_i^\alpha(x,z) - \sum_{j=1}^{N}\psi_j^S(z_j,1)[1-q_j^\alpha(x,z)]\}g(x,z)dxdz \geq 0. \tag{42}$$

Define an α*-mechanism as follows. An α-mechanism is an α*-mechanism only if (i) $G(\alpha^*) = 0$ for some $\alpha^* \in [0, 1]$ or (ii) $G(0) > 0$.

THEOREM 2. Let (F, H) describe a trading problem for which an α*-mechanism exists. Suppose that the distributions (F, H) have the properties that, for every buyer i and seller j, $\bar{p}_i^{\alpha^*}(x_i)$ and $\bar{q}_j^{\alpha^*}(z_j)$ are nondecreasing over the intervals $[a_i, b_i]$ and $[c_j, d_j]$ respectively. The α*-trading mechanism $(p^{\alpha^*}, q^{\alpha^*}, r^{\alpha^*}, s^{\alpha^*})$ is ex ante efficient, individually rational, and incentive compatible for (F, H). Its expected gains from trade are positive.

An Example. The number of buyers equals the number of sellers (N = M) and all traders' reservation values are identically and uniformly distributed on the unit interval. Properties of the α^*-mechanism as the number of buyers and sellers varies are:

N=M	α^*	$\alpha^*/(1+\alpha^*)$	$1/\alpha^*$	Efficiency	Inefficiency
1	.3333	.2500	3.00	.8436	.1564
2	.2256	.1841	4.43	.9437	.0563
3	.1603	.1382	6.24	.9733	.0267
4	.1225	.1091	8.17	.9847	.0153
6	.0827	.0764	12.09	.9931	.0069
8	.0622	.0586	16.08	.9961	.0039
10	.0499	.0475	20.04	.9975	.0025
12	.0416	.0399	24.04	.9983	.0017

Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders

by

Lawrence R. Glosten*
Northwestern University

and

Paul R. Milgrom
Yale University

Abstract

We examine properties of bid and ask prices quoted by a risk-neutral, zero-profit specialist facing heterogeneously informed traders. The presence of traders with superior information leads to a positive bid-ask spread. The resulting transaction prices convey information and the expected average spread during a trading period is bounded by a number which is proportional to the square root of the number of arrivals. An increase in insider activity, an increase in the quality of their information or an increase in the elasticity of liquidity trader demand or supply leads to larger spreads. A bid-ask spread implies a divergence between observed returns and realizable returns. Observed returns are approximately realizable returns plus what uninformed anticipate losing to informed.

*Lawrence R. Glosten
J.L. Kellogg Graduate School of Managmeent
Northwestern University
Evanston, Illinois 60201

THE REVELATION OF INFORMATION IN STRATEGIC MARKET GAMES

J. Geanokopolos
P. Dubey
M. Shubik

Yale University

We develop here a simple and workable alternative to rational expectations equilibrium in which the transmission of information into and out of prices can be explicitly modeled. In this formulation the process of price formation is represented by the Nash equilibrium of a multi-period, playable game. The difficulty, we suggest, with the notion of rational expectation equilibrium (R.E.E.) is that it provides an explanation for how prices in an economy can convey information, but not for how the information gets into the prices in the first place. Furthermore, the difference between our strategic equilibrium and R.E.E. can be illustrated best by the differing interpretations they provide for the efficient markets of hypothesis.

Consider an economy in which agents have different levels of information concerning exogenous random states of nature $s \in S$. In planning their actions, what account do agents take of what others might know about the relative profitability of their opportunities? Does the economy eventually behave, Hayek wondered, as if all the collective information were held in a single agent's hands?

In a R.E.E. each trader is assumed to know the equilibrium price function $p: S \to \Delta^{L-1}$, and to condition his demand on his own information I^n and the particular price realization $\bar{p} = p_s$. It is easy to show that a R.E.E. is individually rational in the sense that no trader can end up with a state contingent commodity bundle which is (ex ante) worse than his initial endowment. The naive competitive equilibrium, in which traders use prices to calculate their budget constraints without inferring any information from them, does not satisfy this basic rationality principle. In the special case where there is no uncertainty, then of course R.E.E. and the usual Walrasion equilibria coincide. Finally, when the state space S is finite, then generically R.E.E. exist and are fully revealing, in the sense that every agent can infer from prices alone the combined knowledge I* of all the other traders. Thus R.E.E. generalizes Walrasion equilibrium to economies with differing levels of information and answers both of the questions of the last paragraph.

But how can it be, we ask, that agents take into account the prevailing price levels $\bar{p}$ in choosing their actions when it is precisely those actions which create the prices? Without an explanation of how prices are formed there can be no intrinsic reason that the prices do not reveal information that no agent originallypossesses. Indeed in a fully revealing R.E.E. no agent has an incentive to expend energy to learn something someone else already knows, for he will be able to infer it for free from prices.

We present a two-period model with an explicit process for the flow of information via prices. In the first period agents act on the basis of their privately held information in submitting bids and supplies for each of the goods. These actions produce prices (the quotients of total bids and total offers for each good) and allocations. In the second period agents act again, but this time they augment their own information with inferences they make from the previous period's prices.

In this model the Nash equilibrium allocations satisfy the individual rationality property-an agent can always choose not to trade. If there is no uncertainty (and a sufficient amount of the money good), then with a continuum of traders the Nash equilibria are identical to the Walvasion equilibria. Our main proposition shows that when the state space if finite, there are, generically a finite number of equilibria. In all of these equilibria the first period prices are fully revealing, so that in the <u>second</u> period agents are acting as if they possessed the combined knowledge of all. Nevertheless, agents who begin with superior information typically profit from the advantage, so that they have an incentive to gather it.

We conclude by pointing out that the bid-offer model we chose as the mechanism of our model could be replaced by many other similar mechanisms, including submitting price-contingent demands (limit orders) with a fixed number of prices, without affecting our results. In fact we show that no continuous mechanism can implement rational expectations equilibria generically for finite state space economics.

ABSTRACT OF

A SEQUENTIAL SIGNALLING MODEL OF CONVERTIBLE DEBT CALL POLICY

by

Milton Harris and Artur Raviv*

February 1984

In this paper we attempt to resolve two puzzles concerning convertible debt calls. The first is that although it has been shown by Ingersoll (1977a) that conversion of these bonds should optimally be forced as soon as this is feasible, actual calls are significantly delayed relative to this prescription [Ingersoll (1977b)]. The second is that common stock returns are significantly negative around the announcement of the call of a convertible debt issue [Mikkelson (1981)]. This observation is also inconsistent with Ingersoll's (1977a) model. Our purpose is to rationalize managers' observed call decisions and the market's reaction to them in a framework in which managers act in the interest of their shareholders and investors' reactions are rational. This is accomplished by a model in which managers periodically receive private information regarding the firm's prospects. Each period the manager must decide whether or not to force conversion of the outstanding convertible debt. The market tries to infer the private information of managers by observing their call decisions. This results in stock value adjustments. Managers take these effects into account along with their private information in making their decisions. The mutual consistency between optimal managerial decisions and optimal market responses is guaranteed by the

*Department of Finance, J. L. Kellogg Graduate School of Management, Northwestern University, Evanston, IL 60201.

use of the sequential equilibrium concept [Kreps and Wilson (1982)].

We show that there is an equilibrium in which managers truthfully signal their private information by calling the convertible bonds if and only if their current information is unfavorable. Intuitively, firms receiving favorable information have lower costs of foregoing an opportunity to force conversion since for such firms it is more likely that conversion will take place voluntarily later on. The benefits of not forcing conversion, which are independent of private information, consist of an increase in current stock value due to investors' belief that such behavior signals favorable information.

The signalling equilibrium exhibits behavior which is consistent with the delay in forcing conversion documented by Ingersoll (1977b). It also explains the negative stock return associated with call announcements since in our equilibrium these are taken by investors to be signals of unfavorable information. In addition to explaining these puzzles, several other implications follow from the signalling equilibrium. First, we show that stock price will rise each time the firm foregoes an opportunity to force conversion. Mikkelson's (1981) evidence suggests that this phenomenon may indeed occur. Second, the longer a firm delays conversion, the better its subsequent performance (in a stochastic dominance sense). Finally, convertible bond prices also increase when forced conversion is foregone and decrease in the event of forced conversion. These are potentially testable implications of our model.

REFERENCES

Ingersoll, J. E. Jr., "A Contingent Claims Valuation of Convertible Securities," Journal of Financial Economics, 4, 1977a, 289-322.

Ingersoll, J., "An Examination of Corporate Call Policies on Convertible

Securities," Journal of Finance, 32, 2, May 1977b, 463-478.

Kreps, D. M. and R. Wilson, "Sequential Equilibria," Econometrica, 50, 4, July 1982, 863-894.

Mikkelson, W. H., "Convertible Calls and Security Returns," Journal of Financial Economics, 9, 1981, 237-264.

A THEORY OF EXIT IN OLIGOPOLY

by

Drew Fudenberg and Jean Tirole

January 1984

University of California, Berkeley, and CERAS, Ecole Nationale des Ponts et Chausees, respectively.

We develop a model in which exit occurs as a group of firms is reduced to those which will remain active in the long run. Each firm enters the market knowing its own costs, but not those of its opponents. As tim goes on, weaker firms leave, and so each firm becomes increasingly pessimistic about the strength of its remaining opponents. The time of exit is the only strategic variable, and there are only two firms, so that our model is a "war of attrition." In contrast to the classic war of attrition in which "fighting" (i.e. both players staying in) is always costly, we assume that with small but-positive probability each firm's costs are low enough that it is profitable as a duopolist. This assumption results in a unique equilibrium, in contrast to previous results.

Price and Advertising Signals of the Quality of New Experience Goods

Paul Milgrom* and John Roberts †

* Yale University † Stanford University

In this paper we present a game-theoretic model in which two (or more) observable choice variables are used in equilibrium to signal for a single, one-dimensional characteristic, the value of which is private information.

The model is cast in terms of a firm which has just introduced a new product whose quality is either high (H) or low (L). The actual, realized quality is known to the firm but is not known by potential customers. Instead, customers have prior beliefs about quality and learn its actual level only through experience with the product. With such "experience goods", it is in the customers' interest to become better informed about quality before purchase and in the interests of the firm, if it is producing high quality, to make this fact known. However, a low-quality producer would also like to be taken as producing high quality, since this gives it access to greater demand. In these circumstances, for example, if the H firm simply claims it is high quality and is believed, then an L will copy it and make the same claims. Thus, if the H firm is to indicate credibly that it is high quality, it must find some act that it can take that an L will be unwilling to mimic. This act would then be a signal for quality.

Some years ago, P. Nelson suggested that the volume of advertising could be a credible signal of quality. His idea was that a high-quality product would attract more repeat sales than a low quality one. Thus, an H firm would be willing to expend more than would an L—on advertising or whatever—to attract initial sales that would provide the basis for repeat sales. Further, if customers come to associate high advertising levels with high quality, then such ads would in fact bring in more initial sales and so the expenditure and the customers'beliefs would be mutually justified.

It is particularly worth noting that these ads need carry no direct information about the product's quality. Indeed, as suggested above, any advertised claims about quality are too freely mimicked to be credible. Further, in the absence of the informational asymmetry, there would be no advertising.

Nelson did not present a formal analysis based on his ideas. Our model fills this gap while incorporating a major element that was missing in his story, viz., that price might also signal quality. This would occur, for example, if costs differed with quality and markups were correlated with quality in the same direction as costs. If prices and qualities are so correlated, of course, then customers can infer quality without paying any attention to ad volume, and it is no longer obvious why the firm should advertise.

A key element of our analysis is that we are able (generically) to obtain a unique signalling equilibrium. Typically, signalling models have a continuum of equilibria. This indeterminacy arises from the latitude that exists in specifying how the signal recipients interpret values of the

*Yale University †Stanford University

signal that ought not to have been observed in equilibrium. Bayes Rule determines their beliefs conditional on observing a value of the signal that would occur with positive probability under the specified strategy for the signaller, but provides no guidance when an event with prior probability zero is observed. Yet these beliefs are crucial determinants of behavior. Because behavior generated by these beliefs off the putative equilibrium path is a crucial determinant of what is in fact an equilibrium, the usual notions of equilibrium thus place few restrictions on behavior and thus permit a multiplicity of equilibria.

Our approach to reducing the set of equilibria involves the sequential elimination of weakly dominated strategies and of strategies that cannot be played with positive probability in a sequential equilibrium. This leads us to specify a locus of points in price-advertising space such that an L would not find it worthwhile to select a (P,A) pair on or beyond this locus, even if by so doing it would initially be thought to be an H. Given this, we specify that choices on or beyond the locus are taken as surely indicating an H. Choices from inside the locus are ones that fail to signal credibly that the firm is producing high quality, since an L would gladly make such a choice if it led to its being thought to be an H. In fact, we specify that customers take such a choice as surely indicating an L.

Given this specification, an H that does decide to signal its quality faces the problem of maximizing its profits (conditioned on its being thought an H) subject to its choice being on or beyond the specified locus (so that it will be thought an H).

The constraint set in this problem is highly non-convex. However, if the full information price-advertising choice for an H, $(P_H, \emptyset)$, does not satisfy the constraint (so that price alone does not credibly signal quality) and if the second partial with respect to price of the profit function for a high quality firm that is viewed as such is strictly more negative than that for a low quality firm that is initially thought to be an H, then a positive level of advertising is used in equilibrium. As well, price will typically differ between L and H firms and the price of the high quality firm will also differ from its full-information value, P_H.

The analysis is illustrated in the accompanying diagrams. In them, P_L denotes the price that would be charged by a low-quality producer that was recognized as such. The locus discussed above is drawn heavily. The other curve in each diagram is a level curve for the H firm's profits. The solution is $(P^*_L, A^*_L) = (P_L, \emptyset)$, since the L is recognized as such in equilibrium, and $(P^*_H, A^*_H) = (P^*, A^*)$. Note that we may have P^* greater or less than P_H, depending on the relationship of P_H and $P^{\emptyset}$, the price that maximizes the profits of an L which is initially thought to be an H.

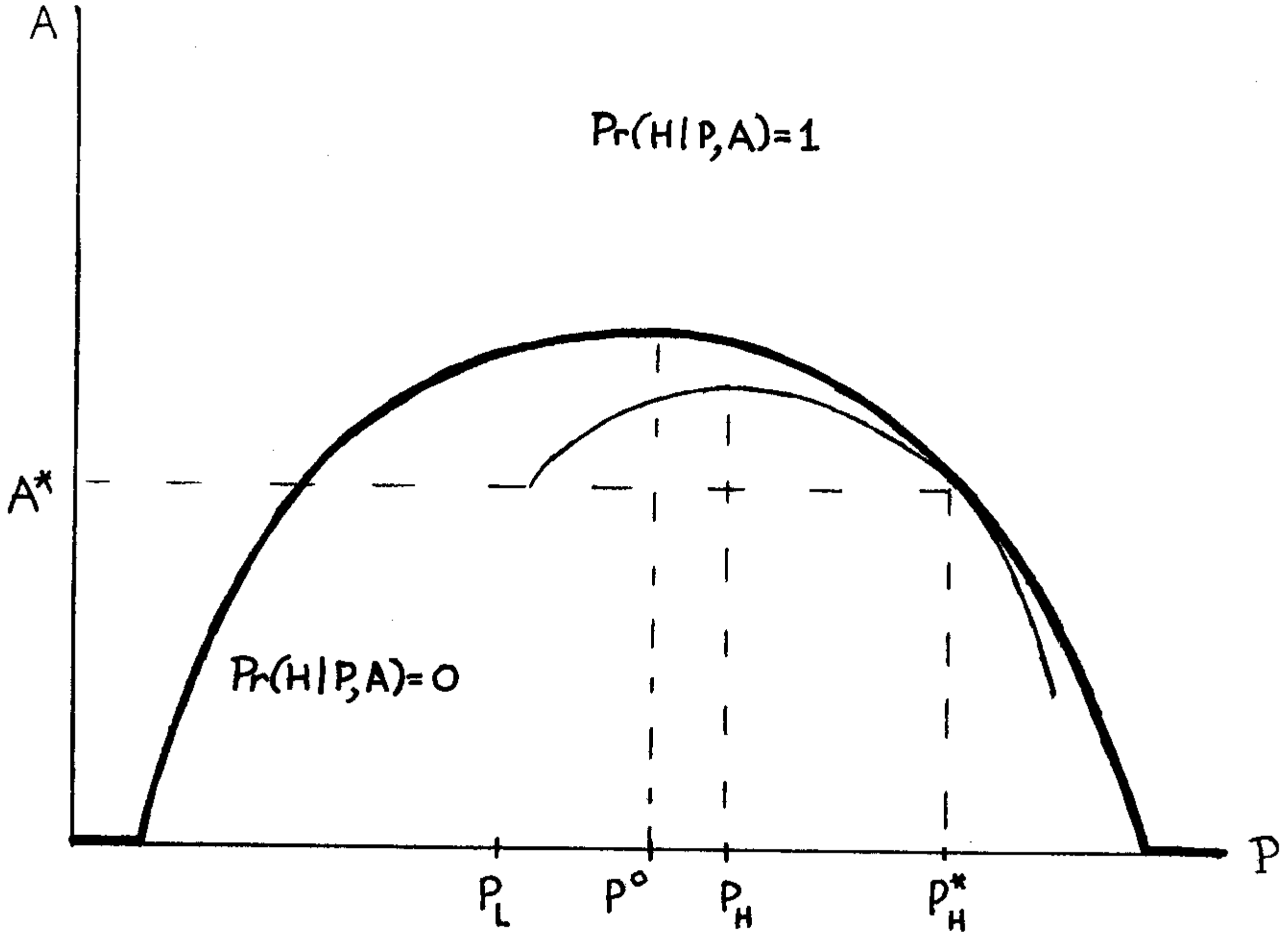
A
Pr(H|P,A)=1
A*
Pr(H|P,A)=0
P
P_L
P^o
P_H
P_H^*

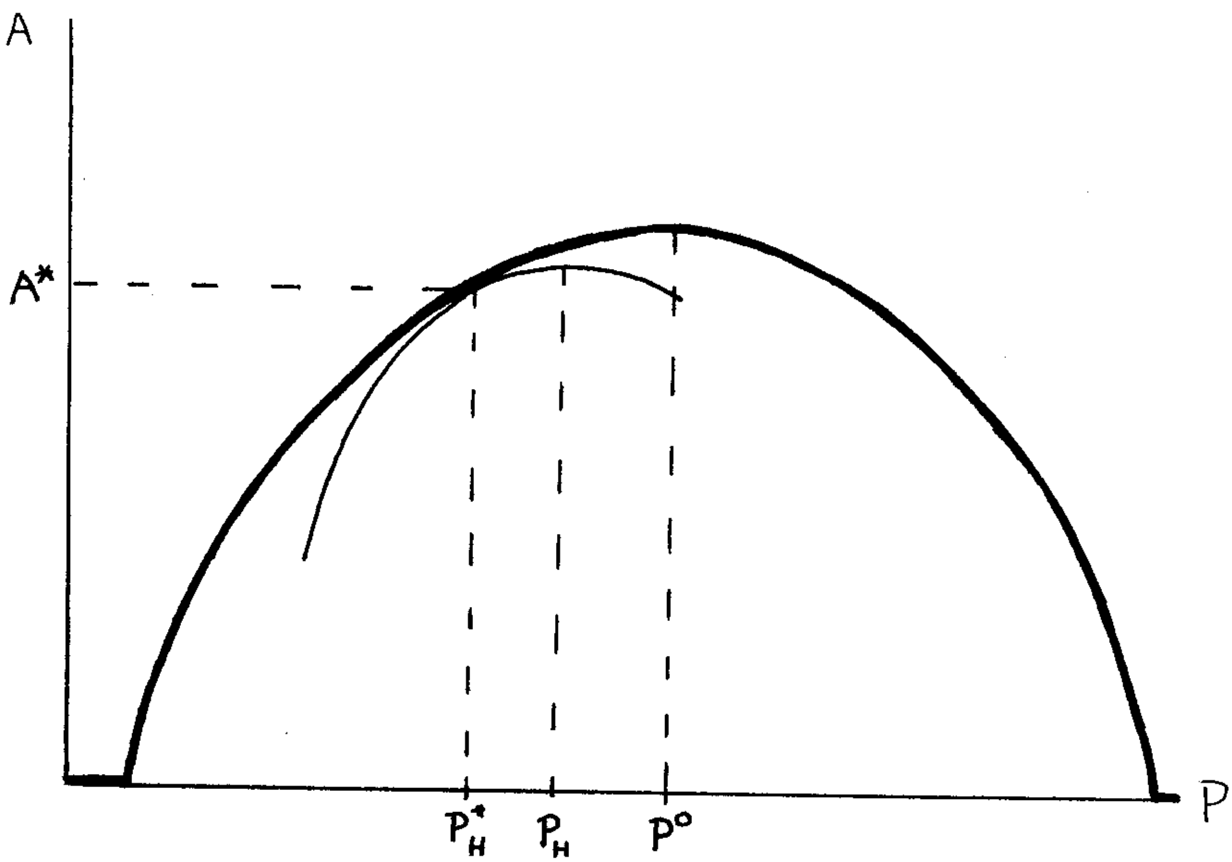
A
A*
P
P_H^*
P_H
P^o

THE ARCHITECTURE OF ECONOMIC SYSTEMS:

HIERARCHIES AND POLYARCHIES

Joseph E. Stiglitz
Princeton University

There is a widespread belief that the internal organization of an economic system has an important effect on its performance. Yet, there is very little in traditional economic analysis which investigates such a relationship. In this paper, we present some new ways of looking at economic systems. We motivate our discussion in the context of economic systems, but it has implications for the internal organization of large corporations as well.

The thesis of this paper is that economic systems behave fundamentally differently under different forms of organization, and that central to an understanding of the performance of an economic system is an understanding of its architecture. The architecture (like that of a computer or electrical system) describes, among other things, how the constituent decision making units are arranged together in a system (i.e., who makes which decisions) and who conveys what information to whom. Our attempt is to relate the architecture of alternative economic systems, along with their attendant rules for meting out rewards and punishments, to their performance.

The axiom of human behavior which plays a basic role in our analysis is that all decision makers make errors of judgment. For concreteness we focus on simple decisions which involve accepting or rejecting certain projects. Individuals (or the constituents of economic systems) make these decisions based on the information available to them. In any event, because of the errors in judgment, some projects which get accepted should have been rejected, and some projects which are rejected should have been accepted. Using the analogy from the classical theory of statistical inference, these errors correspond to Type-II and Type-I errors.

The typology of economic systems on which we focus in this paper arises from the differences in how individuals are organized together in a system. We think of a polyarchy as a system in which there are independent (and possibly competing) sources of decision making. In contrast, a hierarchy is visualized as a system in which the decision making authority is more concentrated. In Section I, we present simple (polar) models of these two systems.

The most important consequence of how the individuals are arranged together is that the aggregation of errors is different in different economic systems. The aggregation of errors, in turn, determines the performance of a system. For example, in a market economy, if one firm rejects a profitable idea (say, for a new product), then there is a possibility that some other firm might accept it. In contrast, if a single agency makes such decisions and this agency rejects the idea, then the idea must remain unused. The same, however, is also true for those ideas which are unprofitable. As a result, one would expect a greater incidence of Type-II error in a polyarchy, and a greater incidence of Type-I error in a hierarchy.

It should be apparent, however, that the overall performance of a system (for example, its profit level) will depend not only on its architecture, but also on the mix of projects that is available to its decision makers, and on the nature of errors that the decision making entails.

Initially, in Section II, we assume that individuals make similar errors in the two systems, and that the nature of errors is exogenously specified.

Moreover, the mix of projects available to the two systems is also identical. The performance of economic systems is thus attributable primarily to what we have called their architecture.

We then examine, in Section III, how the portfolio of projects available to an economic system is influenced by its architecture. This represents more of a general equilibrium view: the differences in the architecture will affect the chances that different types of projects have of being accepted or rejected. This will, in turn, influence the incentives of those who conceptualize and invent projects, and will thus affect the kinds of projects which are invented. For example, one would expect that the inventors would attempt to generate those projects which are more likely to be accepted.

The above analysis takes the errors in judgment (that is, the probabilities of good projects being rejected and those of bad projects being selected by evaluators) as exogenous, as well as identical in the two systems. In Section IV, we analyze an endogenous determination of these errors. This we do by determining rational screening rules for project acceptance, where the constituents in the economic systems take into account whatever information is available to them.

In this paper, we analyze only one, albeit an important one, aspect of the architecture of economic systems. Some other important aspects are sketched out in the concluding section.

ABSTRACT

INFORMATION, PLANNING AND CONTROL IN HIERARCHIES

by

John Geanakoplos and Paul Milgrom

Yale University

We study a resource allocation problem in which individual managers have internalized the organization's objectives but are limited in the rate at which they can process information for decision making. Under those circumstances, delegating parts of the decision to each of several managers results in a greater volume of relevant information being processed, but it also results in poorer coordination of decisions. That loss of coordination is alleviated by the creation of a hierarchy.

Managers in our model use information to reduce costs of production. In the production of cost-saving decisions, managerial attention is a nonhomogeneous input: since different decisions are delegated to different managers, it matters which manager attends to which information. Moreover, if the information a manager can process is proportional to the time he devotes, then there are generally increasing returns to any one manager's time. This, by itself, may explain the long hours worked by some managers and the more than proportionate pay that such managers sometimes receive.

For a quadratic example, we characterize the product of management, and show how the organization of tasks in a hierarchy influences that product. It pays to group productive units in ways that ease the information processing of the unit managers and their superiors and to arrange decisions so that the cost impact of residual uncertainties is reduced.

Adding a "normal-normal" Bayesian information processing technology to the quadratic example, we find that improved information processing

leads eventually to larger spans of control. Also, increases in prior knowledge about an organization's technological parameters diminish the marginal product of managerial "talent". These facts suggest that older, more stable industries should have relatively low ratios of managerial wages to value added, or should employ less able managers at the lowest levels, or both.

The preceding assumptions are still not sufficient to imply that there is a bound on the size of a hierarchy, or that it pays to assign the ablest managers to the top slots. To generate those conclusions, we use the assumptions (i) that managers in a hierarchy cannot begin making decisions until the higher level managers have reached their decisions, (ii) that the time devoted to decision making at each level is chosen optimally, and (iii) that the planning manager is kept fully employed in the planning role. It then follows that the marginal value of a manager's time rises with his level in the hierarchy.

The rising marginal value of time can account for the use at higher levels in the hierarchy of abler, more diligent managers and of time-saving services and perquisites. It may be possible to understand chauffeured limousines, company jets, executive assistants, and sophisticated decision support systems all from this point of view. Similarly, the size of an optimal hierarchy may be limited, since the opportunity cost of the time spent in planning at the top level rises with the size of the organization, being roughly proportional to the number of managers at the lowest level.