Federal Reserve Bank of Minneapolis
Research Department Staff Report 226

December 1996

Implementation Theory with Incomplete Information*

Beth Allen

Federal Reserve Bank of Minneapolis
and University of Minnesota

ABSTRACT

This paper surveys implementation theory when players have incomplete or asymmetric information, especially in economic environments. After the basic problem is introduced, the theory of implementation is summarized. Some coalitional considerations for implementation problems are discussed. For economies with asymmetric information, cooperative games based on incentive compatibility constraints or Bayesian incentive compatible mechanisms are derived and examined.

*This paper is based on the author's NATO Advanced Study Institute lecture at Stony Brook and prepared for inclusion in the volume entitled *Cooperation: Game Theoretic Approaches,* edited by Sergiu Hart and Andreu Mas-Colell, to be published by Springer-Verlag, Heidelberg, Germany. This work was supported by the National Science Foundation (through grant SBR93-09854), the Federal Reserve Bank of Minneapolis, and the Curtis Carlson Chair in Economics at the University of Minnesota. Mr. Xiaodong Huang of the University of Minnesota was the notetaker for the lecture on which this paper is based. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

In this paper, we examine part of the literature regarding implementation under incomplete or asymmetric information. Implementation includes not only the classical social choice problem of characterizing the functions or sets of allocations or outcomes that can be obtained as the result of a group decision—in this case, with truthful Bayes-Nash equilibrium in a direct mechanism—but also the extension of these group decisions to encompass cooperative game-theoretic solution concepts (rather than exclusively non-cooperative equilibria) and the inclusion of incentive compatibility or mechanism considerations into cooperative games derived from economies with asymmetric or incomplete information.

After introducing the basic incentive compatibility problem, we proceed to examine implementation theory proper when there is incomplete information, with particular emphasis on Jackson's (1991) necessary and sufficient conditions for Bayes-Nash implementation. After some remarks concerning the possibilities for considering coalitional behavior in the implementation problem with incomplete information, we redirect our attention to economies with incentive compatibility constraints and the games they generate. Finally, we briefly consider a game-theoretic model of how agents could cooperatively select a Bayesian incentive compatible mechanism.

2. The Basic Problem

Consider a simple prototype implementation problem with asymmetric information. Suppose that the uncertainty is summarized by a set $S = \{s_1, s_2, s_3\}$ of states of the world or signals correlated with states. Let A be a set of possible actions. The problem is to pick a mapping from states to actions optimally.

Suppose further that there are two individuals, one of whom (distinguished by the subscript I for "informed") knows the state $s \in S$ while the other (denoted by the subscript U for "uninformed") does not know anything about which state $s \in S$ has occurred, although the set S, an objective probability on it, the utility functions, and the information structure are all common knowledge. The preferences of these two individuals are given by state-dependent utilities $u_I : A \times S \to I\!R$ and $u_U : A \times S \to I\!R$.

Incentive compatibility here means that $a : S \to A$ satisfies $u_I\big(a(s); s\big) \geq u_I\big(a(s'); s\big)$ for all $s, s' \in S$. This is also sometimes termed the self-selection constraint. It means that the informed agent is willing to reveal truthfully the state of the world, because in all states $s \in S$, the utility of $a(s)$ given s is never less than the utility the informed agent could obtain by stating $s' \in S$ and thus receiving $a(s')$ when the true state is $s \in S$. There is no incentive compatibility constraint for the second agent because his lack of information is common knowledge.

For the special case in which $A \subset I\!R$ and u_I is strictly monotone on A for each $s \in S$, incentive compatibility requires $a(s_1) = a(s_2) = a(s_3)$. The informed agent cannot be forced to reveal information truthfully if doing so would lead to this agent receiving less "money" than he or she could obtain in some other state of the world.

Changing the model to give the uninformed agent partial information so that he can distinguish $\{s_1\}$ from the event $\{s_2, s_3\}$ can alter the results. However, whether this partial information is verifiable—whether it can be confirmed by some third party who can act as a referee if necessary—matters greatly. If the information is verifiable, the partially uninformed agent can force an allocation in state s_1 which may be either better or worse for the fully informed agent than what the fully informed agent received in s_2 or s_3 [which still must satisfy $a(s_2) = a(s_3)$ in the strictly monotone one-dimensional example], and similarly for the informed agent. If $\{s_1\}$ versus $\{s_2, s_3\}$ is not verifiable, then we must add incentive compatibility constraints for the partially uninformed agent in order to force him to reveal correctly whether he believes that the true state lies in $\{s_1\}$ or $\{s_2, s_3\}$.

For convenience, one sometimes imposes excess incentive compatibility when there is asymmetric information. Doing so may decrease welfare, but it sometimes doesn't change the qualitative properties of the solution. Obviously, the advantage is to simplify notation by, for instance, treating informed and uninformed agents symmetrically by giving them all the same incentive compatibility constraint that properly applies only to the informed agents (when their identities are common knowledge). This procedure may be correct if the less informed agents cannot be distinguished and if all agents are permitted to announce the state of the world; in this case, the uninformed agents could always, for instance, announce the state $\bar{s}$ if $u_U\big(a(\bar{s}); s\big) \geq u_U\big(a(s'); s\big)$ for all $s, s' \in S$.

In the terminology of noncooperative game theory, incentive compatibility says that telling the truth is a Nash equilibrium in the game with strategies consisting of announcements about states of the world and payoffs defined by utilities evaluated at the proposed $a : S \to A$ mapping for the true state of the world. Implementation basically means that an allocation can be obtained as a truth-telling Nash equilibrium; this idea will be made more precise later.

An introduction to this literature can be found in d'Aspremont and Gérard-Varet (1979, 1982), Myerson (1991), and Postlewaite and Schmeidler (1987). Note, however, that I shall not attempt to give a complete reference list or even a historical summary of this topic.

3. Implementation with Asymmetric Information

A fundamental result in implementation theory is the revelation principle, which *roughly* states that anything which is incentive compatible (and hence implementable) can be implemented as a truth-telling (Nash) equilibrium of a direct mechanism, where a direct mechanism is a noncooperative game in which players' strategies consist of complete announcements of what they know about their "type" (i.e., preferences). The extension to incomplete information frameworks is due to Rosenthal (1978), Myerson (1979), and Harris and Townsend (1981); for a discussion, see the textbook by Myerson (1991, pp. 260-261) or the survey paper by Postlewaite and Schmeidler (1987).

Bayes-Nash Revelation Principle. *With incomplete information, if an allocation function can be obtained as a Bayes-Nash equilibrium (of some mechanism or some* *communication game), then it can be implemented with truthful equilibrium strategies in a direct mechanism.*

An important insight is the importance of an informational condition, known as publicly predictable information (PPI) or nonexclusivity of information (NEI). The assumption states that no player has information which is not at least as coarse as the pooled information of all other players. In symbols, if we let the sub-σ-field $\mathcal{G}_i$ denote the information of player $i \in N$, where all of the $\mathcal{G}_i$ are sub-σ-fields of a given σ-field $\mathcal{T}$ of measurable events, publicly predictable information precisely requires that for all $i \in N$, $\mathcal{G}_i \subseteq \sigma\big(\bigcup_{j \neq i} \mathcal{G}_j\big)$. This means that all of the other players, acting together, can always detect lies by any individual. The significance of publicly predictable information is that it permits the use of "forcing contracts" or mechanisms in which an extremely bad outcome arises whenever a single player tells a lie. If the messages sent by players are inconsistent, the mechanism assigns the worst possible outcome so that any unilateral lie looks extremely risky; hence, truth must be a Nash equilibrium. The condition was discovered by—in alphabetical order—Blume and Easley (1990), Palfrey and Srivastava (1987), and Postlewaite and Schmeidler (1986); see also the discussion by Postlewaite and Schmeidler (1987).

An important research topic was the elucidation of necessary conditions and sufficient conditions for Bayes-Nash implementation. This work has resulted in a huge literature, including the articles by Blume and Easley (1990), Palfrey and Srivastava (1987), and Postlewaite and Schmeidler (1986) mentioned above. Contributions by Palfrey and Srivastava (1989) and Jackson (1991) are especially relevant here; Jackson's (1991) result for economic environments will be discussed in detail in the following section because he does obtain a set of conditions which are both necessary and sufficient. Further literature includes articles by Matsushima (1988, 1991) and Palfrey and Srivastava (1986, 1991). See also the recent survey by Palfrey and Srivastava (1992) and the background material on games with communication due to Forges (1986) and Myerson (1986).

Palfrey (1992) focuses on the problem of multiple equilibria for Bayes-Nash implementation. Mechanisms—even those direct mechanisms for which truth telling is a Nash equilibrium—typically exhibit many Nash equilibria. Therefore, the value of the revelation principle may be limited in the sense that some allocation could well be implementable as the unique equilibrium of some mechanism while being only one of a plethora of equilibria of direct mechanisms for which the given allocation arises as the truthful equilibrium. The notion of full implementation addresses this issue, as full implementability of an allocation means that it is implementable as the unique Bayes-Nash equilibrium of some suitable mechanism.

Ledyard (1986) expounds a critique of the concept of implementation. Using the mild hypotheses of strictly positive prior probabilities and monotonically increasing transformations of utilities, he points out that any undominated outcome can be rationalized as a Bayes-Nash equilibrium of some game. Of course, this means that Bayes-Nash implementation doesn't lead to interesting restrictions unless one either tightens the requirements of the definition of implementation (for instance, by requiring full implementation), restricts the class of allowable games, or insists on some refinement of Bayes-Nash equilibrium.

4. Jackson's Article

Jackson (1991) provides a set of necessary and sufficient conditions for Bayes-Nash implementation with or without the hypothesis of publicly predictable information. Previous work using the PPI assumption found two necessary conditions for Bayes-Nash implementability: incentive compatibility (also called self-selection) and a Bayesian analogue of Maskin's (1977) monotonicity condition. [See Postlewaite and Schmeidler (1986), Palfrey and Srivastava (1987), and Blume and Easley (1990).] In attempting to find a converse result, Jackson (1991) adds a closure condition [which was somewhat implicit in the Palfrey and Srivastava (1987) and Postlewaite and Schmeidler (1986) assumptions] that one can always "patch together" allocation functions at common points in players' information partitions. [The operation is reminiscent of Savage's (1954) treatment of personal probability and expected utility.] Subject to technicalities, incentive compatibility, Bayesian monotonicity, and closure are together necessary and sufficient for Bayes-Nash implementation.

For Jackson's (1991) theorem, we consider an exchange economy with at least three traders and strictly monotone utilities for every trader in every state of the world. The formulation could allow for public goods and externalities. To fix notation, let $N = \{1, \ldots, n\}$ be the set of economic agents and let $\preceq_i$ denote trader i's preference relation. I shall define and explain terminology after stating the result. To simplify, I restrict attention to economic environments; see Jackson (1991) for extensions to more general situations.

Theorem. *A social choice set is implementable if and only if there exists a social choice set $\widehat{F}$ which is equivalent to F such that $\widehat{F}$ satisfies (IC), (BM), and (C).*

Definition 1. A *social choice set* F is a subset of the set of all social choice functions. In symbols, if $S = S_1 \times \ldots \times S_n$, where for $i \in N$, S_i is the finite information set of player i, and if A denotes the set of all feasible acts, which are assumed to be independent of elements in the set S (i.e., let A be the set of state-dependent allocations that are resource-feasible, given traders' initial endowments $e_i \in I\!R^\ell_+$ for $i \in N$, so that $A = \{(\widetilde{x}_i : S \to I\!R^\ell_+)_{i \in N} |$ for all $s \in S$, $\sum\limits_{i \in N} \widetilde{x}_i(s) = \sum\limits_{i \in N} e_i\}$), then F is a subset of $X = \{x | x : S \to A\}$.

We say that two social choice sets are *equivalent* if they are equal almost surely. Consequently, we need only work with those social choice sets that are defined on some convenient subset of S of full measure. If every $s = (s_1, \ldots, s_n) \in S$ occurs with strictly positive probability, no two distinct social choice sets can be equivalent; in this case, the theorem reduces to the statement that F is implementable if and only if it satisfies (IC), (BM), and (C).

Definition 2. A social choice set F satisfies *condition* (IC) if for all $i \in N$, all

$x \in F$, all $s \in S$, and all $t_i \in S_i$, $x(s) \succeq_i (s_i)\ x(s)_{i(}, t_i)$, where $\succeq_i (s_i)$ denotes trader i's preference relation when his information set is $s_i \in S_i$.

Definition 3. A social choice set F satisfies *condition* (C) if for all common knowledge partitions $\{S', S''\}$ of S and all $x, y \in F$, there is $z \in F$ such that $z(s) = x(s)$ if $s \in S'$ and $z(s) = y(s)$ if $s \in S''$.

The closure condition is needed because equilibria of mechanisms can similarly be patched together based on common knowledge events. If a mechanism has two Bayes-Nash equilibria—call them x and y—then it must also have a third equilibrium, z, defined by doing x on part of S and y on the other part of S, providing that S can be divided into two or more pieces that are common knowledge.

Definition 4. F satisfies *condition* (BM) if, whenever $x \in F$ and $\alpha = (\alpha_1, \ldots, \alpha_n)$ is a *deception*, where $\alpha_i : S_i \to S_i$ for all $i \in N$ and $x \circ \alpha$ denotes the social choice function with outcomes $x(\alpha(s)) = x(\alpha_1(s_1), \ldots, \alpha_n(s_n))$ for all $s = (s_1, \ldots, s_n) \in S$, and whenever there is no social choice function in F which is equivalent to $x \circ \alpha$, then there exists $i \in N$, $s_i \in S_i$ and $y \in X$ such that $(y \circ \alpha) \succ_i (s_i)\ (x \circ \alpha)$ and $x \succeq_i (t_i)\ y \circ \alpha_i(s_i)$ for all $t_i \in S_i$, where $\big(y \circ (\alpha_i(s_i))\big)(s) = y\big(s)_{i(}, \alpha_i(s_i)\big)$ for all $s \in S$.

An interpretation of the Bayesian monotonicity condition is as follows (ignoring equivalence): If a mechanism implements F and if $x \in F$, then there is an equilibrium σ (of the game defined by the mechanism) which yields x. If agents use deception α, they obtain $x \circ \alpha$. If there does not exist a social welfare function in F which is equivalent to $x \circ \alpha$, then $\sigma \circ \alpha$ cannot be an equilibrium. Bayesian monotonicity ensures that, in fact, $\sigma \circ \alpha$ isn't an equilibrium. The idea is that agent i uses y to signal that α is being played; this makes trader i happier. The second condition says that player i cannot gain by falsely accusing others of deception.

5. Cooperative Implementation

By definition, implementation is a noncooperative concept; it requires allocations to arise as (truthful) Nash equilibria. Perhaps the most straightforward way to include the consideration of coalitional behavior is to replace Nash equilibrium with strong equilibrium. A disadvantage of this approach is that strong equilibria may not exist in general noncooperative games, whereas there are always Nash equilibria, at least in mixed strategies under fairly general technical conditions. A more radical strategy is to examine the possibilities for attaining outcomes as some cooperative solution in a game. In this case, the precise application of incentive compatibility constraints is unclear. Should one worry about incentives to lie within a coalition that is cheating? Are blocking allocations required to be incentive compatible? Such considerations seem to have the flavor of bargaining sets (i.e., objections versus counterobjections) or coalition proofness.

Incentive compatibility can be incorporated into cooperative games on three levels. First, one can find some solution set and ask whether it satisfies incentive compatibility or, as a weaker alternative, at least contains some outcome satisfying incentive compatibility. This is the approach taken by Krasa and Yannelis (1994) and Koutsougeris and Yannelis (1992). Secondly, one can require incentive compatibility only in the definition of feasible actions for the grand coalition. This approach implicitly appears in the second-best efficiency considerations for the planner in the literature on incentives and mechanism design. Finally, one can consistently require incentive compatibility for the definition of feasible agreements for all coalitions. This strongest use of incentive compatibility treats all coalitions symmetrically but possesses the disadvantage of possibly leading to games that violate some of the standard properties one expects. This tack is followed in Allen (1991, 1992, 1993, 1994).

A further factor which complicates the analysis is that games without transferable utility are more appropriate when incentive considerations are present. To summarize the worth of a coalition by a single number—as is done in the definition of coöperative games with transferable utility (or *TU* games)—suggests that members of the coalition share a single objective function. Yet, if these players were indeed a team, they would necessarily be willing to share their information fully and honestly in order to better maximize the total payoff accruing to the coalition. This contradicts the spirit of incentive compatibility, which hypothesizes that players will hide information or will lie to further their own goals.

Finally, one can ask whether participation or individual rationality constraints should be imposed. Requiring that all players be willing to play the game is natural for some mechanism problems, as it is a weaker rationality requirement than Bayesian incentive compatibility. On the other hand, in a cooperative context, most solution concepts are automatically—by definition—individually rational, although out-of-equilibrium behaviors such as blocking and objecting may not always be individually rational compared to nonparticipation. Moreover, *ex ante* and *ex post* individual rationality are distinct concepts. The latter restricts risk sharing so that its imposition can prevent efficient outcomes such as those obtainable with fair insurance contracts.

6. Economies with Asymmetric Information

Consider a pure exchange economy with agent set $N = \{1, \ldots, n\}$ in which Ω is a finite set. To simplify, assume that every state of the world occurs with strictly positive probability and that these probabilities are common knowledge. Agents' information either consists of partitions on Ω or can be specified by signals $\mathbf{s}_i : \Omega \to S_i$, where each S_i is also assumed to be a finite set. Write $S = \Pi_{i \in N} S_i$ and $\mathbf{s} = (\mathbf{s}_1, \ldots, \mathbf{s}_n)$. Consumption sets are $\mathbb{R}^\ell_+$ and initial endowments are $e_i \in \mathbb{R}^\ell_+$ for $i \in N$. Endowments are assumed not to depend on Ω or S in order to guarantee that initial endowment vectors are incentive compatible and hence that there exist incentive compatible feasible allocations. Preferences are specified by state-dependent cardinal utilities $u_i : \mathbb{R}^\ell_+ \times \Omega \to \mathbb{R}$ where, for every $i \in N$ and every $\omega \in \Omega$, $u_i(\cdot; \omega) : \mathbb{R}^\ell_+ \to \mathbb{R}$ is continuous, strictly monotone, and strictly concave.

The classical incentive compatibility constraints are given by the restrictions that allocations $x_i : \Omega \to \mathbb{R}^\ell_+$ must satisfy, for all $i \in N$ and all $\omega \in \Omega$, $u_i\big(x(\omega); \omega\big) \geq$

$u_i\big(x_i(\omega');\omega\big)$ for all $\omega' \in \Omega$. Note that these constraints apply to every player regardless of the coalition to which he belongs. They are written in "overkill" fashion, as if each player were able to distinguish all states rather than in a form that reflects the player's individual information (which could depend on his coalition). Think of these incentive compatibility constraints as restrictions on the state-dependent consumption set of each agent.

Alternatively, for a framework in which traders receive signals about the state of the world, Bayesian incentive compatibility requires

$$\sum_{\omega\in\Omega} u_i\big(x_i(\mathrm{s}(\omega));\omega\big)\mu_i(\omega|s_i) \geq$$
$$\sum_{\omega\in\Omega} u_i\big(x_i(s_i', \mathrm{s}_{)i(}(\omega));\omega\big)\mu_i(\omega|s_i)$$

for all $s_i \in S_i$, all $s_i' \in S_i$, and all $i \in N$, where the allocation $x_i : \Omega \to I\!R^{\ell}_+$ must be measurable with respect to the signals $\mathrm{s}(\cdot) = \big(\mathrm{s}_1(\cdot), \ldots, \mathrm{s}_n(\cdot)\big)$, and $\mu_i(\omega|s_i)$ denotes player i's posterior probability of $\omega \in \Omega$, given that he or she has observed signal $s_i \in S_i$.

7. Incentives with Asymmetric Information

The study of cooperative solution concepts for economies with incentive considerations has focused primarily on the core, although the value has also been examined. One approach that has proved useful is to analyze the cooperative games with nontransferable utility that are generated by (exchange) economies with incentive compatibility constraints. Thus, one defines the cooperative games $V : 2^N \to I\!R^n$ with nontransferable utility (or NTU games) by $V(\emptyset) = I\!R^n$ and for $T \subseteq N$ with $T \neq \emptyset$, $V(T) = \big\{(w_1, \ldots, w_n) \in I\!R^n |$ for $i \in T$, there exists $x_i : \Omega \to I\!R^{\ell}_+$ such that, for all fully informed $i \in T$ and all ω, $\omega' \in \Omega$, $u_i\big(x_i(\omega);\omega\big) \geq u_i\big(x_i(\omega');\omega\big)$, where $\sum_{i\in T} x_i(\omega) = \sum_{i\in T} e_i$ for all $\omega \in \Omega$ and $w_i \leq \sum_{\omega\in\Omega} u_i\big(x_i(\omega);\omega\big)\mu(\omega)$ for all $i \in T\big\}$, where $\mu(\omega)$ is the probability of state ω and agents in N are assumed to be either fully informed (i.e., their information partitions on Ω precisely equal 2^{Ω}) or completely uninformed (i.e., their information partitions on Ω are the trivial partitions $\{\Omega\}$). One can modify the game to take careful account of players' partial information or to use the Bayesian version of incentive compatibility constraints.

The incentive compatible core was first introduced by Boyd and Prescott (1986) in a model of financial intermediation with risk neutrality. They demonstrate nonemptiness of the core by showing that certain systems of linear inequalities can be solved. Berliant (1992) and Marimon (1989) also examine the incentive compatible core for particular economic problems—those involving taxation and adverse selection. Allen (1991, 1994) follows the approach outlined above of deriving *NTU* games from economies with (classical or Bayesian) incentive compatibility constraints and finds that the game need not be balanced and can, in fact, have an empty core. For economies with asymmetric information, Koutsougeris and Yannelis (1992) define core allocations and check whether they are incentive compatible.

For the value, Allen (1992) derives the games from economies with (classical or Bayesian) incentive compatibility and shows that the value is well defined. Krasa and Yannelis (1994) focus on the private information value and ask whether the fine, coarse, and private information values satisfy incentive compatibility.

8. Mechanisms with Asymmetric Information

Instead of adding incentive compatibility constraints to the definition of the games derived from economies with asymmetric information, one can incorporate Bayesian incentive compatible mechanisms into the definition of these games. This approach builds on the work of Harsanyi (1967-68) on noncooperative games with incomplete information and its use by Myerson (1984) to model cooperative games with incomplete information.

Allen (1993) proposes a game containing both cooperative and noncooperative phases in which the feasible outcomes are taken to be Bayesian-Nash equilibrium outcomes of Bayesian incentive compatible direct mechanisms. Formally, the entire model is assumed to be common knowledge and, in the first strategic phase, players cooperatively pick a Bayesian incentive compatible mechanism. The choice of a mechanism is a binding agreement; the commitment is made *ex ante*. Then, after agents learn their types, the noncooperative game defined by the chosen mechanism is played. Traders send messages (about their types, since we can restrict ourselves to direct mechanisms by the Bayes-Nash revelation principle), which lead to an outcome according to the mechanism. The equilibrium concept used in the noncooperative (mechanism) game phase is Bayesian-Nash equilibrium, which (by the revelation principle) can be taken to be truthful. Somewhat more formally, the game given by $V(S) = \{(w_1, \ldots, w_n) \in \mathbb{R}^n |$ there exists a randomized direct Bayesian incentive compatible mechanism λ, and there is a (truthful) Bayesian-Nash noncooperative equilibrium σ for λ such that, if $i \in S$, i's payoff in λ under σ is at least as great as $w_i\}$. The use of incentive compatible mechanisms in cooperative economic contexts is also studied by Ichiishi and Idzik (1992), Page (1992), and Rosenmüller (1990).

Acknowledgments. This work was supported by the National Science Foundation (through grant SBR93-09854), the Federal Reserve Bank of Minneapolis, and the Curtis Carlson Chair in Economics at the University of Minnesota. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Mr. Xiaodong Huang of the University of Minnesota was the notetaker for the lecture on which this paper is based.

References

Allen, Beth (1991) Incentives in market games with asymmetric information: The core. Center for Analytic Research in Economics and the Social Sciences Working Paper #91-38, Department of Economics, University of Pennsylvania. Also distributed as CORE Discussion Paper No. 9221, Center for Operations Research and Econometrics, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.

Allen, Beth (1992) Incentives in market games with asymmetric information: The value. Center for Analytic Research in Economics and the Social Sciences Working Paper #92-11, Department of Economics, University of Pennsylvania.

Allen, Beth (1993) Supermechanisms. Mimeo, Department of Economics, University of Minnesota.

Allen, Beth (1994) Incentives in market games with asymmetric information: Approximate *NTU* cores in large economies. In: Barnett, W.A., Moulin, H., Salles, M., Schofield, N.J. (eds.) *Social Choice, Welfare, and Ethics. Proceedings of the Eighth International Symposium in Economic Theory and Econometrics, Caen, 1992.* International Symposia in Economic Theory and Econometrics, Cambridge University Press, 263-276.

d'Aspremont, Claude and Gérard-Varet, Louis-André (1979) Incentives and incomplete information. *Journal of Public Economics,* **11**, 25-45.

d'Aspremont, C. and Gérard-Varet, L.-A. (1982) Bayesian incentive compatible beliefs. *Journal of Mathematical Economics,* **10**, 83-103.

Berliant, Marcus (1992) On income taxation and the core. *Journal of Economic Theory,* **56**, 121-141.

Blume, Lawrence and Easley, David (1990) Implementation of Walrasian expectations equilibria. *Journal of Economic Theory,* **51**, 207-227.

Boyd, John H. and Prescott, Edward C. (1986) Financial intermediary-coalitions. *Jour-*

nal of Economic Theory, **38**, 211-232.

Forges, Françoise M. (1986) An approach to communication equilibria. *Econometrica,* **54**, 1375-1385.

Harris, Milton and Townsend, Robert M. (1981) Resource allocation under asymmetric information. *Econometrica,* **49**, 33-64.

Harsanyi, John C. (1967-68) Games with incomplete information played by "Bayesian" players. *Management Science,* **14**, 159-182, 320-334, 486-502.

Ichiishi, Tatsuro and Idzik, Adam (1992) Bayesian cooperative choice of strategies. Mimeo, Department of Economics, Ohio State University.

Jackson, Matthew O. (1991) Bayesian implementation. *Econometrica,* **59**, 461-477.

Koutsougeris, Leonidas and Yannelis, Nicholas C. (1992) Incentive compatibility and information superiority of the core of an economy with differential information. Mimeo, Department of Economics, University of Illinois at Champaign-Urbana.

Krasa, Stefan and Yannelis, Nicholas C. (1994) The value allocation of an economy with differential information. *Econometrica,* **62**, 881-900.

Ledyard, John O. (1986) The scope of the hypothesis of Bayesian equilibrium. *Journal of Economic Theory,* **39**, 59-82.

Marimon, Ramon (1989) The core of private information economies. UAB/IAE Discussion Papers W.P. 131.90, Universitat Autònoma de Barcelona.

Maskin, Eric (1977) Nash equilibria and welfare optimality. Mimeo, Department of Economics, Massachusetts Institute of Technology.

Matsushima, Hitoshi (1988) A new approach to the implementation problem. *Journal of Economic Theory,* **45**, 128-144.

Matsushima, Hitoshi (1991) Incentive compatible mechanisms with full transferability. *Journal of Economic Theory,* **54**, 198-205.

Myerson, Roger B. (1979) Incentive compatibility and the bargaining problem. *Econometrica,* **47**, 61-74.

Myerson, Roger B. (1984) Cooperative games with imcomplete [*sic*] information. *International Journal of Game Theory,* **13**, 69-96.

Myerson, Roger B. (1986) Multistage games with communication. *Econometrica,* **54**, 323-358.

Myerson, Roger B. (1991) *Game Theory: Analysis of Conflict.* Harvard University Press.

Page, Frank H. Jr. (1992) Bayesian incentive compatible mechanisms. Mimeo, Department of Finance, University of Alabama at Tuscaloosa.

Palfrey, Thomas R. (1992) Implementation and Bayesian equilibrium: The multiple equilibrium problem in mechanism design. In: Laffont, Jean-Jacques (ed.) *Advances in Economic Theory, Volume 1. Proceedings of the Sixth World Congress of the Econometric Society, Barcelona, 1990.* Econometric Society Monograph Series, Cambridge University Press, 283-323.

Palfrey, Thomas R. and Srivastava, Sanjay (1986) Private information in large economies. *Journal of Economic Theory,* **39**, 34-58.

Palfrey, Thomas R. and Srivastava, Sanjay (1987) On Bayesian implementable allocations. *Review of Economic Studies,* **54**, 193-208.

Palfrey, Thomas R. and Srivastava, Sanjay (1989) Implementation with incomplete information in exchange economies. *Econometrica,* **57**, 115-134.

Palfrey, Thomas R. and Srivastava, Sanjay (1991), Efficient trading mechanisms with pre-play communication. *Journal of Economic Theory,* **55**, 17-40.

Palfrey Thomas R. and Srivastava, Sanjay (1992) *Bayesian Implementation.* Fundamentals of Pure and Applied Economics, Harwood Academic Publishers.

Postlewaite, Andrew and Schmeidler, David (1986) Implementation in differential information economies. *Journal of Economic Theory,* **39**, 14-33.

Postlewaite, Andrew and Schmeidler, David (1987) Differential information and strategic behavior in economic environments: A general equilibrium approach. In:

Groves, T., Radner, R., and Reiter, S. (eds.) *Information, Incentives, and Economic Mechanisms: Essays in Honor of Leonid Hurwicz.* University of Minnesota Press, 330-348.

Rosenmüller, Joachim (1990) Fee games: $(N)TU$-games with incomplete information. Institut für Mathematische Wirtschaftsforschung Working Paper Nr. 190, Universität Bielefeld.

Rosenthal, Robert W. (1978) Arbitration of two-party disputes under uncertainty. *Review of Economic Studies,* **45**, 595-604.

Savage, L.J. (1954) *The Foundations of Statistics.* Wiley, New York.