Econometrica, Vol. 60, No. 6 (November, 1992), 1273–1292

ASYMPTOTIC EFFICIENCY IN LARGE EXCHANGE ECONOMIES WITH ASYMMETRIC INFORMATION

BY FARUK GUL AND ANDREW POSTLEWAITE[1]

We provide conditions on an exchange economy with asymmetric information that guarantee that when the economy is replicated sufficiently often, there will be an allocation which is incentive compatible, individually rational, and nearly efficient. The main theorem covers both the case in which aggregate uncertainty remains when the economy is replicated and the case in which replication eliminates aggregate uncertainty. In addition, we demonstrate how our theorem does or does not apply to standard asymmetric information problems such as the buyer's bid double auction problem, Akerlof's lemons problem, and insurance with asymmetric information.

KEYWORDS: Implementation, incentive compatibility, incomplete information, general equilibrium, noncooperative games, large economies, rational expectations equilibrium.

1. INTRODUCTION

THERE HAS BEEN IN RECENT YEARS voluminous research showing the various ways that asymmetric information among agents in an economy can preclude the attainment of a (first best) Pareto efficient outcome. It is plausible that every economic situation involves some asymmetry of information: every agent probably knows something about his utility function or production technology that is not known to all other agents. It then follows that in every economic problem, there may be inefficiency due to the asymmetric information.

Despite the seeming ubiquity of this asymmetric information induced inefficiency, many economists believe that in many circumstances competitive markets generate efficient, or nearly efficient, outcomes. Although seldom spelled out, this belief is based on a notion that while there may be asymmetry in agents' information, it is relatively unimportant for the problem at hand because any single agent has only a small amount of information not known by the other agents.

It is an attractive idea that there should be a concept of an agent's being informationally small, and when agents are informationally small, the inefficiency due to asymmetric information is small. There are two immediate difficulties with the notion, however. First, it may be that the aggregate effect of many agents, each of whom is informationally small, prevents efficient market performance. A second difficulty with the notion is that the measure of informational smallness isn't obvious. One of the first of the papers showing that

[1] This paper is a major revision of an earlier paper with the same title. This research was partially supported by the National Science Foundation. Gul thanks the Alfred P. Sloan Foundation for financial assistance. We would also like to acknowledge research assistance from the Graduate School of Business, Stanford University. We gratefully acknowledge helpful conversations with Richard Kihlstrom, Andreu Mas-Colell, Hugo Sonnenschein, and Bob Wilson. We also acknowledge very helpful comments from two referees and a co-editor. We particularly want to thank Hideo Suehiro for carefully reading and pointing out numerous errors in an earlier draft.

markets may perform poorly in the presence of asymmetric information, Akerlof's (1970) lemons paper, has a large number of agents each of whom has private information only about the car that he has for possible sale. By many criteria, it might seem that each agent is informationally small here and yet (what is plausibly) the market outcome is as far as possible from the outcome that would have arisen had the asymmetry not existed.

Various models aimed at describing the effect of asymmetric information on the performance of markets can be distinguished according to how they model the market mechanism and how they define a market outcome. Often, models of financial markets begin with a particular extensive form game, designed specifically to mimic a given trading institution, and investigate the efficiency properties of that game. This approach has the advantage that it enables conclusions regarding the actual performance of the given institution. However, these conclusions are often very sensitive to minor modelling decisions (implicit in the choice of the extensive form game) about which we are often less than certain. We avoid this difficulty by choosing not to model a particular mechanism but rather to take the essential feature of all market mechanisms, voluntary participation both in trade and in the sharing of information (that is, individual rationality and incentive compatibility), and use the revelation principle to investigate the restriction imposed on efficiency by voluntary participation.

As a result, our theorems identify classes of situations in which markets might (rather than will) be efficient. Insisting on voluntary participation (in the above sense) also distinguishes our approach from those that use rational expectations equilibrium to define market outcomes since rational expectations equilibrium prices often appear to convey information that agents would not willingly provide given the consequences of its availability. The typical rationale for using the rational expectations equilibrium concept is presumably that agents are informationally small and therefore market prices would only be slightly affected by the information of any single agent and competition will presumably lead agents to reveal their information. One of our main objectives is to get a better sense of what it means to be informationally small. Thus, we are led to use a framework in which incentive constraints are modelled explicitly.

We stress that our aim in this paper is *not* to argue that large numbers of agents will *always* eliminate the inefficiencies caused by asymmetry of information among agents. We will provide conditions that ensure that large numbers of agents will, in fact, eliminate those inefficiencies. These conditions will be satisfied in some problems with asymmetric information but will fail to hold in others. We will show in Section 3 that two frequently studied problems involving asymmetric information, the market for lemons studied by Akerlof and the market for insurance with adverse selection, do *not* satisfy the conditions of our theorem.

A sufficient condition for (essentially) eliminating the incentive problem turns out to be the following: the incremental impact of each agent's information (given the information of others) on the demand of every good should be small. Under our assumptions, increasing the number of agents ensures this. We will explain why informational smallness in the above sense is not attained in Akerlof's model and insurance problems with adverse selection even when the number of agents is large.

In Section 2.1 we offer a simple example of the type of structure analyzed in this paper. The example aims to provide insight into our Theorem, the proof of which is somewhat complicated. Section 2.2 contains our model and results. We leave to Sections 3 and 4 a discussion of our results and their relation to other papers.

2.1. *Example*

There are two people who live in a duplex, Andy (A), who lives on the top floor, and Bob (B) who lives on the bottom floor. B owns a smoke alarm; the alarm is of more value to A than to B since A lives on the top floor and is more likely to be trapped by a fire than would B. Neither person will benefit from the alarm unless the alarm is in his home. The values to each person depend upon the probability of there being a fire, of course. Suppose it is common knowledge that A does not smoke, but only B knows whether B smokes or not. It is common knowledge that A believes the probability that B smokes to be $1/2$. The possibility that B smokes is exogenous; call the state in which B smokes s and the state in which he does not smoke n. The utilities to A and B from having the smoke alarm and $\$m$ in the states s and n are given as follows:

	state:	
agent	s	n
A	$m + 10$	$m + 3$
B	$m + 8$	$m + 1$

The utility to either person of $\$m$ with no smoke alarm is m.

We can see that regardless of the state, the only efficient outcome is that the alarm be transferred from B to A. If the outcome is to be (ex post) individually rational, A must give B between \$8 and \$10 in the state s and between \$1 and \$3 in state n. If there is a mechanism that achieves an outcome for this problem that is ex post Pareto efficient and ex post individually rational, it must be the outcome of a revelation game in which the agents report their private information truthfully. Since only agent B has private information, this is easy to check. It is easy to see that the best we could do to relax the incentive constraints on agent B while maintaining individual rationality would be to have agent A pay \$8 in state s and \$3 in state n. But it is clear that agent B will not reveal truthfully if he faces these prices, but will announce state s regardless of the true state. Thus there is an incompatibility between incentive compatibility and the achievement of individually rational and Pareto efficient outcomes.

Now suppose we replicate the problem in the following way. Let there be r duplex houses each containing a pair that looks like the pair A and B above. In each duplex the person living on the bottom will own the smoke alarm and will have exclusive information as to whether he smokes or not. Assume these potential smokers will each have probability of 0.5 of smoking and that whether smokers in different duplexes smoke is independent.

How does this r-fold replication change the conflict between efficiency and incentive compatibility? In fact we will see that the replication process asymptotically eliminates the conflict. Consider the following mechanism for reallocating the smoke alarms. All smoke alarms will be sold for \$5 or not sold at all.[2] Every B is to announce whether or not he smokes. Those announcing that they smoke will keep their smoke alarm and not trade. Those announcing that they do not smoke will sell their smoke alarms to the upper level tenants of the buildings inhabited by B's who have announced that they smoke. If there are fewer B's who announce that they smoke than those who announce that they do not smoke, there will be excess supply of alarms; in this case we ration the sellers uniformly by selecting the appropriate number of sellers randomly from the set of B's who have announced that they do not smoke. In the opposite case, that in which there are more B's who say they smoke than say they do not, there will be excess demand for alarms. We uniformly ration the buyers in this case.[3]

It is easy to see that B's will announce truthfully in the face of this scheme. Any B who does not smoke should say that he does not so as to have some possibility of selling his alarm for a price higher than his reservation price. Lying will guarantee that he does not trade, yielding a lower expected utility. Next, a B who does smoke should say he does since the best that can happen is that he sells the alarm for less than his reservation value. Hence the outcome is incentive compatible.

The outcome is trivially individually rational by the construction of the mechanism. The sense in which it is asymptotically efficient is as follows. Every alarm that is transferred goes from a person whose value is \$1 to a person whose value is \$10, for a gain of \$9 or \$4.50/person trading. The inefficiency is that there may be excess supply of alarms (more persons announcing they do not smoke than announcing that they do), or excess demand (fewer smokers than nonsmokers). In the first case, there will be smoke alarms left in the hands of the rationed sellers who value them at \$1 when there are people who value alarms at \$3 without them. In the case of excess demand, there will be unsatisfied buyers who value the alarms at \$10 and smokers who kept their alarms when they valued them at \$8.

But the proportion of the alarms that are in the hands of people who value them less than others without alarms is going to zero. By the law of large numbers, the probability that the actual number of smokers is arbitrarily close to $1/2$ is going to 1. Thus with probability 1, the proportion of the potential gains from trade (in the absence of private information) that are realized by the mechanism is going to 1.

In the next section, we present our model and results. In Section 3 we discuss whether several important problems involving asymmetric information do or do

[2] The price can as well be any price between 3 and 8.

[3] The mechanism used in proving the main theorem does not use rationing but rather taxes every trading agent by a fixed small amount.

not fit into our model. We leave to Section 4 a general discussion of our results and their relation to other papers.

2.2. *The Model*

Let $N=\{1,2,\dots,n\}$ be a set of economic agents. Let $\tilde{\theta},\tilde{t}_1,\tilde{t}_2,\dots,\tilde{t}_n$ be $n+1$ distinct random variables each with the following properties:

(i) Finite support: There exist finite sets $\Theta, T_1, T_2,\dots,T_n$ such that $P(\tilde{\theta}\in\Theta, \tilde{t}_i\in T_i \text{ for } i=1,2,\dots,n)=1$.

(ii) Full range: For any $\theta\in\Theta$ and $t\in\prod_{i=1}^n T_i$ $P(\tilde{\theta}=\theta,\ \tilde{t}=t)>0$.

(iii) Nontriviality: For any θ,θ' with $\theta\neq\theta'$, there exists $t\in\prod_{i=1}^n T_i$ such that $P(t|\theta)\neq P(t|\theta')$.

Henceforth we will use $P(\theta)$, $P(t_i)$ etc. to denote $P(\tilde{\theta}=\theta)$ and $P(\tilde{t}_i=t_i)$. Also note that the probability spacc on which the random variables are defined is suppressed. We will use T to denote $\prod_{i=1}^n T_i$ which is the support of the random vector $(\tilde{t}_1,\tilde{t}_2,\dots,\tilde{t}_n)$. The symbols $t,\bar{t}$ will be used to denote generic elements of T and $t_i, t_i', \hat{t}_i$ will be used to denote generic elements of T_i. Also we will use t_{-i} to denote $(t_1,t_2,\dots,t_{i-1},t_{i+1},\dots,t_n)$ and (t_{-i},t_i') to denote $(t_1,t_2,\dots,t_{i-1},t_i',t_{i+1},\dots,t_n)$.

The consumption set of agent i is $\boldsymbol{R}^l_+$ and $w_i\in\boldsymbol{R}^l_{++}$ is the endowment of agent i. The utility index u_i of agent i is a function u_i: $\boldsymbol{R}^l_+\times T\times\Theta\to\boldsymbol{R}$ where $u_i(\cdot,t,\theta)$ is continuous, strictly concave, increasing, and bounded for every $t\in T$ and $\theta\in\Theta$. The collection $e=((\tilde{t}_i,w_i,u_i)_{i\in N},\tilde{\theta})$ will be called a private information economy PIE. Let $D_i(\cdot,t,\theta)$ denote the demand function of an agent with preferences $u_i(\cdot,t,\theta)$ and endowment w_i. Note that $D_i(p,t,\theta)$ is well-defined for any $p\in \operatorname{Int} S^{l-1}$ (where $\operatorname{Int} S^{l-1}$ is the interior of the $l-1$ dimensional simplex). We will be interested in P.I.E.'s that satisfy the following regularity condition: for all $i\in N$, $\varepsilon>0$, $p\in\operatorname{Int} S^{l-1}$ and $t_i',\hat{t}_i\in T_i$ such that $t_i'\neq\hat{t}_i$, there exists p', t_{-i}, and θ s.t. $D_i(p',(t_{-i},t_i'),\theta)\neq D_i(p',(t_{-i},\hat{t}_i),\theta)$ and $\|p-p'\|<\varepsilon$.

Thus the regularity condition requires that the demands of two different types for an agent should never be identical for all prices in some open ball for every realization of the relevant uncertainty.

An allocation $x=(x_1,x_2,\dots,x_n)$ for the PIE e is a collection of functions x_i such that x_i: $T\to\boldsymbol{R}^l_+$.

An allocation x for the PIE is *incentive compatible* (IC) if

$$U_i(x|t_i)\geqslant U_i(x,\hat{t}_i|t_i)\quad\text{for all}\quad i\in N,\ t_i,\hat{t}_i\in T_i\quad\text{where}$$

$$U_i(x|t_i)=\sum_\theta\sum_t U_i(x(t),t,\theta)P(\theta,t|t_i)\quad\text{and}$$

$$U_i(x,\hat{t}_i|t_i)=\sum_\theta\sum_{t_{-i}} U_i\big(x(t_{-i},\hat{t}_i),(t_{-i},t_i),\theta\big)P(\theta,t_{-i}|t_i).$$

An allocation x is said to be:

Feasible (F) if $\sum_i(x_i(t)-w_i)=0$ for all $t\in T$.

Ex Post Individually Rational[4] (XIR) if

$$\sum_\theta u_i(x_i(t),t,\theta)P(\theta|t) \geqslant \sum_\theta u_i(w_i,t,\theta)P(\theta|t)$$

for all $i \in N$ and $t \in T$.

Ex Post Efficient (XE) if (i) x is feasible; (ii) for any allocation y,

$$\sum_\theta u_i(y_i(t),t,\theta)P(\theta|t) \geqslant \sum_\theta u_i(x_i(t),t,\theta)P(\theta|t)$$

for all $i \in N$ and $t \in T$ and for some $\bar{t} \in T$ and

$$j \in N \sum_\theta u_j(y_j(\bar{t}),\bar{t},\theta)P(\theta|\bar{t}) > \sum_\theta u_j(x_j(\bar{t}),\bar{t},\theta)P(\theta|\bar{t})$$

implies y is not feasible.

Ex Post ε-efficient ($X_\varepsilon E$) if (i) x is feasible; (ii) there exists $E \subset T$ such that $P(t \in E) > 1 - \varepsilon$ and for any allocation y,

$$\sum_\theta u_i(y_i(t),t,\theta)P(\theta|t) \geqslant \sum_\theta u_i(x_i(t),t,\theta)P(\theta|t) + \varepsilon$$

for all $i \in N$ and $t \in E$ implies y is not feasible.

We will be concerned with r-fold replicas of a given PIE. For any $e = \{(\tilde{t}_i, w_i, u_i)_{i\in N}, \tilde{\theta}\}$ consider a PIE $e^r = \{(\tilde{t}_{is}, w_{is}, u_{is})_{(i,s)\in N\times R}, \tilde{\theta}\}$ where $R = \{1,2,\dots,r\}$. e^r is said to be an r-fold or r replica of e if:

(i) $w_{is} = w_i$ for all $s \in R$;

(ii) the joint distribution of $\tilde{\theta}, \tilde{t}_{is}, \tilde{t}_{2s}, \dots, \tilde{t}_{ns}$ is the same as the joint distribution of $(\tilde{\theta}, \tilde{t}_1, \tilde{t}_2, \dots, \tilde{t}_n)$ for all $s \in R$;

(iii) for any $\theta \in \Theta$, $i, j \in N$, $s, \hat{s} \in R$ such that $s \neq \hat{s}$, $t'_{is} \in T_i$, $t'_{j\hat{s}} \in T_j$, $P(\tilde{t}_{is} = t'_{is}$ and $\tilde{t}_{j\hat{s}} = t'_{j\hat{s}}|\theta) = P(\tilde{t}_{is} = t'_{is}|\theta)P(\tilde{t}_{j\hat{s}} = t'_{j\hat{s}}|\theta)$;

(iv) $u_{is}(v, t^r, \theta) = u_i(v, t^r_s, \theta)$ for all $v \in \boldsymbol{R}^l_+$, $i \in N$, $s \in R$, and $t^r \in T^r$ where $t^r = (t^r_1, t^r_2, \dots, t^r_r)$ and $t^r_{\hat{s}} \in T$ for all $s \in R$.

Thus an r-fold replica of e contains r "copies" of each agent $i \in N$. A priori each copy of an agent i is identical, i.e., has the same endowment and "same" preferences. However, the preferences of agents depend on θ and the information of other agents in their "cohort" (i.e., $s \in R$). Furthermore, the realization of type profiles across cohorts is independent given the true value of $\tilde{\theta}$. Thus as r increases each agent is becoming "small" in the economy both in terms of endowment and information. Note that for large r an agent may have a large amount of private information about the preferences of a small fraction of the economy (i.e., his own cohort) and/or a small amount of private information regarding the preferences of everyone (through his information about $\tilde{\theta}$).

THEOREM: *Let $e = \{(\tilde{t}_i, w_i, u_i)_{i\in N}, \tilde{\theta}\}$ be a private information economy satisfying the regularity condition. Then for any $\varepsilon > 0$ there exists $\bar{r}$ such that for all $r > \bar{r}$ there is an allocation x^r for the PIE e^r (i.e., the r-fold replica of e) which satisfies IC, IR, and $X_\varepsilon E$.*

[4] By ex post we mean conditional upon the realization t. Expectations are still taken over Θ since by assumption, it will never be observed.

While we leave the proof to the Appendix, we will provide a brief outline of the logic. The agents announce their types, and the most likely θ conditional upon the vector of announcements is determined. Given this θ, we consider the artificial nonrandom economy that has a distribution of agents' characteristics the same as the distribution on T conditional upon θ and determine the net trades each possible realization t of T (that is, each possible cohort) would receive in a competitive equilibrium allocation for this artificial economy. By the law of large numbers, the actual distribution of agents' characteristics in a particular realization (given θ) of a replicated economy will be arbitrarily close to the distribution of characteristics in this artificial economy if there are sufficiently many replications.

Suppose then, that the bundle each agent receives as a function of the vector of announcements is the bundle he would receive in the competitive equilibrium of the artificial economy. If all agents are announcing truthfully, this bundle is trivially individually rational since it is what the agent would receive in a competitive equilibrium bundle in the artificial economy. Would any agent have any incentive to report other than his true type? There are two possibilities. First, a different announcement might change the estimate of the most likely θ, and hence the competitive equilibrium allocation used to generate his bundle. However, as the number of replicas gets large, the probability that this could happen goes to zero. The second possibility is that the agent's misrepresentation does not alter the most likely θ, but changes the bundle that the agent gets. But the envy-free property of competitive equilibrium guarantees that such a change can never improve an agent's welfare. In fact, we utilize the regularity assumption to construct bundles that make each agent receive a strictly less desirable outcome with positive probability by misrepresenting. Thus as r increases the second effect dominates the first and hence the unique optimal strategy is to report truthfully.

There is one complication with the argument above; we know that with probability nearly one the distribution of agents' characteristics in the realization of the replica economy will be close to the distribution of characteristics in the artificial economy, but not identical. This means that the allocation in which each cohort gets the vector that it would have gotten in the competitive equilibrium allocation for the artificial economy is not feasible. To handle this problem, the bundles from the competitive allocation for the artificial economy are adjusted in a way that ensures feasibility without upsetting either incentive compatibility or individual rationality. The approximate ex post efficiency of the outcome follows from the fact that each agent receives a bundle that is a competitive equilibrium bundle for an economy not much different from the actual realization of the random economy.

3. APPLICATIONS AND LIMITATIONS OF THE ANALYSIS

3.1. *Informational Smallness and Multilateral Bargaining*

The model as described in the previous section differs sufficiently from other models of asymmetric information that it is worth commenting on its interpretation and application to familiar economic problems involving asymmetric information. A PIE is composed of a set of agents, each of whom has preferences that depend on his information which is represented by the realization t_i, the information of others, t_{-i}, and some unobservable parameter $\tilde{\theta}$. In spite of the fact that $\tilde{\theta}$ is not observed, the agents' information causes them to revise their prior beliefs about $\tilde{\theta}$. Note that the unobservability of $\tilde{\theta}$ necessitates that allocations do not depend on it.

One of the main objectives of this paper is to formalize the sense in which agents need to be informationally small in order to render nearly efficient ex post allocations attainable. The theorem of this paper and its proof suggest that the following definition of informational smallness is suitable for this purpose.

An agent is informationally small in the market for some commodity v if, for most realizations of the relevant uncertainty, the incremental effect of the agent's information (given the information of other agents) on the demand for v is a small proportion of the aggregate endowment of v. The agent is informationally small if she is informationally small in every market v. Notice that it is not necessary for an agent to have a small amount of information; the incremental information of an agent given the information of everyone else should be small. Consider the ways an agent may fail to be informationally small: first, even if the agent has private information only about her own preferences, in general she will not be informationally small if she owns, or is the principle demander of, a large fraction of the aggregate endowment; second, an agent will fail to be small if her incremental information affects the preferences of a large fraction of the population. The replica structure of e^r avoids the first difficulty by ensuring that every agent's endowment becomes an arbitrarily small proportion of the aggregate endowment and that there are increasingly many other agents with the same preferences. To see how the second difficulty is avoided, note that the incremental effect of an agent's information is significant only on her own cohort (i.e., a small fraction of the population). Furthermore, while the private information of the agent on the aggregate parameter of the economy θ may be significant, the incremental effect of this information is also small given the realization of types in the other cohorts.

With this definition of informational smallness, we can see how the proof works. Take any "state" w. Let p and $x(w)$ be Walrasian equilibria for the complete information economy associated with w. The information smallness requirement ensures that we can, by wasting a small proportion of the aggregate endowment, ensure that each agent prefers the allocation she receives in this economy to the demand of any of her other types t_i'. The discreteness (which follows from the finiteness of Θ and T) enables us to do this for every relevant state individually, thus yielding an IC allocation. Ex post individual rationality is also ensured by taking a Walrasian allocation. The simplest example of an economic problem to which our model can be applied is the case of a "private values" economy. Consider the case in which $\tilde{\theta}$ is degenerate, i.e., there exists θ' such that $\text{Prob}(\tilde{\theta} = \theta') = 1$, and suppose that $n = 1$. Thus, the economy to be replicated consists of a single individual whose preferences depend on his private information, t_i. The r-fold replica of this economy consists of r independently drawn agents; each of these agents knows his own preferences and the distribution from which other agents' preferences were drawn. The theorem states that for this private values case, if the number of such independently drawn agents is sufficiently large, there are random allocations for the PIE that are incentive compatible, individually rational and nearly (ex post) Pareto efficient.

It is useful to compare the above example with the work of Gresik and Satterthwaite (1989).[5] Gresik and Satterthwaite considered an environment in which there is a good which might be bought or sold and a number of buyers and sellers, each of whom knows his own reservation price for the good but only the distribution from which other agents' reservation prices were drawn. In this model, Gresik and Satterthwaite prove that the expected inefficiency of an optimal trading mechanism—that is, the mechanism that maximizes expected surplus from trading—goes to zero as the number of traders gets large. Further, they calculate the rate at which the convergence occurs.

To embed this problem in our model, one would consider an initial economy with $\tilde{\theta}$ again taking on a unique value and n equal to the number of sellers plus the number of buyers. Our replication process then mimics the way in which they let the number of agents go to infinity. The difference between the models is that while Gresik and Satterthwaite restrict themselves to risk neutral agents and unitary demand, our model allows general preferences and commodities. Of course Gresik and Satterthwaite's main contribution is computing the rate of convergence while our result deals only with convergence to an ex post efficient allocation.

In the private values case described above, it is clear that when the number of replicas of the basic economy is large, aggregate uncertainty is small since we are taking a large number of independent draws from a given distribution. The fact that aggregate uncertainty is disappearing is *not* fundamental to the asymptotic elimination of the incentive problem. To see this, notice that we can easily represent aggregate uncertainty by identifying different distributions of agents' preferences with different θ's. Thus, even for a large economy there is uncertainty about the distribution of agents' preferences (and hence the set of Pareto efficient allocations). The proposition states that, as in the no aggregate uncertainty case, the incentive problem asymptotically disappears. This discussion points out one further way in which our model is more general than that of Gresik and Satterthwaite: they assume independently drawn reservation prices while in our model, the case in which $\tilde{\theta}$ is nondegenerate allows for a particular nonindependence of the reservation values.

[5] In closely related work Satterthwaite and Williams (1989) analyze the rate of convergence to ex post efficiency of the buyer's bid double auction mechanism for the environments described below. Wilson (1987) shows that the equilibria of the buyer's bid double auction converges to an efficient outcome.

The private values example illustrates a problem for which our proposition assures that replicating an economy will asymptotically eliminate the conflict between efficiency and incentive compatibility. As we emphasized in the introduction, our aim is not to argue that large numbers will always eliminate the conflict between incentive compatibility and Pareto efficiency. Our model and the replication process we employ are not appropriate for some economic problems of interest.

3.2. *Akerlof's Lemons Problem and Insurance with Asymmetric Information*

Consider the application of our model to Akerlof's lemons problem. To see more clearly the problem that arises here, we will return to the example in Section 2.1 and reinterpret it in a manner consistent with the lemons problem. B owns a car and is considering selling it to A. B knows whether he has maintained the car well or not. In either case, A values the car at \$2 more than B, but the actual value depends upon whether or not the car has been maintained.

There is clearly no change from our earlier conclusion that there is a conflict between incentive compatibility and efficiency. Suppose we now replicate the example. In the initial economy $n = 2$, with the first agent being A and the second B. $\tilde{\theta}$ is degenerate, i.e., Θ is a singleton (since it is known that each of a seller's two types has probability .5). $\tilde{t}_1$ is also degenerate since the buyer has no private information; $\tilde{t}_2$ takes on two different values (each with probability $1/2$), representing the seller's private information that the car has been maintained or not. Each agent's utility function depends on the realization of $\tilde{t}_2$. If we replicate this initial economy one time, we have an economy with two A's and two B's. There are now two cars, owned by the two B's; each of the four agents values the cars depending upon t_{21} and t_{22}, the types of the two sellers.

It is easy to verify that replicating the lemons model in this manner is inconsistent with the assumptions of our theorem (since the number of commodities is increasing and the types of agents in one cohort enter the utility functions of agents in other cohorts[6]). Given the intuition outlined above, it is also easy to see why the theorem fails for this particular replication process: The incremental effect of a seller's demand for his own car remains large even as r increases.

One might object to this explanation on the grounds that the definition of a commodity is somewhat ambiguous and that we have arbitrarily chosen a particular definition consistent with our intuitive explanation. Thus one might argue that we could just as well have defined the commodity space as $\boldsymbol{R}_+^2$ (i.e., cars $\times$ money) or $\boldsymbol{R}_+^3$ (i.e., good cars $\times$ bad cars $\times$ money).

[6] Whether it is the increasing number of commodities or the fact that an agent's utility depends upon the types of other agents that is at the heart of the problem is discussed in the last section.

The intuition above rests on the assumption that two objects, x and y, will be considered the same commodity only if they yield the same utility in every state and that they will be considered different commodities only if every agent can distinguish between them (given her information) in every state.[7] The first criterion would rule out the possibility of choosing $\boldsymbol{R}_+^2$ as the commodity space and the second criterion would rule out $\boldsymbol{R}_+^3$.

We will consider next the problem of the provision of insurance in the face of asymmetric information. While there may be more than one way in which the problem can be embedded into our model, we will describe one way to do so that points out the similarity of the problem to that of the lemons problem. Suppose there is a publicly observable event (we will call it an "accident") that affects the value of some asset belonging to an agent. Suppose further that there are two agents, one of whom knows precisely the probability of the accident, while the other knows only the distribution over the possible probabilities, each of which is strictly between zero and one.

The commodity set in this case will be vectors of the consumption good(s) contingent upon the event that the accident has occurred or not. The agent who knows the precise probability will have a type that corresponds to his private information, the probability of accident. Both agents' utility functions will depend on this type, since it determines the probability that they will get the vector of consumption goods contingent upon the accident; as long as the probability of an accident is strictly between zero and one, the agents' preferences will be strictly increasing in these contingent commodity vectors. For this example, Θ is again a singleton; if there were multiple possible distributions over the probabilities of accidents, one could use a nondegenerate $\tilde{\theta}$ to index the distributions.

The above describes an initial economy for the insurance problem. When we replicate the economy, we have four agents. In the standard insurance problem, agents' accidents are typically independent. Thus, the consumption set will now be vectors of consumption goods contingent upon either agent having or not having an accident; that is, the number of commodities has increased as it did in the lemons problem. Also as in that problem, agents utilities depend directly upon the types of agents in cohorts other than their own. Thus for reasons much the same as in the lemons problem, the insurance problem falls outside the scope of our result.

Again, it is worth noting that the incremental effect of each agent's information on the demand for her own insurance does not become small as the number of agents increases. As in the lemons model, it is easily verified that, for the r replica economy, the commodity space must be $\boldsymbol{R}_+^{r+1}$ (insurance contracts for the r agent $\times$ money) in order to satisfy our criterion on the definition of a commodity.

[7] Note that this is indeed a reasonable definition of a commodity when dealing with incentive problems. If the first condition fails, agents will not be indifferent between two different ways of consummating a particular trade. If the second condition fails, agents will not be able to verify if the prescribed trade has indeed taken place.

4. DISCUSSION

1. We assumed that the agents' initial endowments were independent of their types. Among other consequences, this assures that the total endowment in an economy is nonrandom. Since the set of feasible allocations is known, we don't have to confront the possibility of infeasible outcomes. Our theorem can be extended to the case of random endowments by rewriting the proof in terms of net trades provided each agent knows her own endowment given by her type j.

2. As mentioned above, the private values problem in which agents are drawn from a finite set is covered by our theorems. One would certainly expect the conclusion of the theorem to hold in the case in which a sequence of agents is drawn from a distribution over a compact set of agents' characteristics. This involves only an extension from a finite number of possible types to a compact set of possible types. The techniques used in the proofs of our theorems don't work in this case, however. So far, we have been unable to extend our results to include this case. Thus, it is not simply an expositional convenience that we have limited our attention to the finite case.

3. There are two possible types of limiting results for the environment we consider in this paper. A relatively simple result would be that it is possible to design a game in which truthful revelation of an agent's private information (his type) leads to a Pareto efficient allocation and truthful revelation is "nearly" a Bayes equilibrium, where "nearly" means that no agent can achieve more than a small utility gain by misreporting.[8] A simple example of such a game is one that associates the Walrasian equilibrium allocation for the complete information economy described by the agent's announcements. Under quite general assumptions, any agent who misreports his private information can have only a vanishingly small effect on the Walrasian price as the economy gets large. Hence, the utility gain to any agent who misreports will also become vanishingly small as the economy gets large.

But if we are to take seriously the possibility that agents behave strategically, we should assume that they will not content themselves with approximately optimal choices.[9] The problem is potentially very important because while approximately optimal behavior may result in efficient outcomes, this certainly doesn't imply that precisely optimal behavior will result in approximately efficient outcomes. The cumulative effect of many individual agents' adjustments from approximately optimal behavior to optimal behavior, the subsequent adjustments to these adjustments, etc., can be large. There is no reason to expect that there will be a Bayes equilibrium anywhere near an approximate Bayes equilibrium. Thus we are led to the approach in the paper: Design a game in which the Bayesian equilibrium allocations are nearly efficient.

4. In this paper we have taken an approach that is sometimes called weak implementation. We ask only that for a given game, there be *some* equilibrium

[8] This is analogous to the result of Roberts and Postlewaite (1976) for the case of dominant strategy mechanisms for complete information economies.

[9] Or alternatively, we should model explicitly why they do so.

that is nearly efficient when the economy is replicated; there may also be many other equilibria that are not close to being efficient. There are two comments we will make regarding this.

First, even weak implementation answers the question "To what extent does asymmetric information prevent the attainment of an efficient allocation in an economy?" The second comment is that there is a good possibility that substantially more complicated games than that presented in this paper might strongly implement nearly efficient outcomes. That is, there may be extensions of the game presented here that will ensure not only that there is some equilibrium allocation that is nearly efficient, but that *all* equilibrium allocations are. This is a fruitful subject for further research.

5. Since our proof only establishes the existence of some $X_\varepsilon E$, XIR, and IC allocation(s)[10] it is reasonable to inquire if there are other allocations with these properties and, if so, whether there is any particular reason to focus on the allocations that we have constructed. To answer the first part, note that for the argument of our proof to work it is sufficient to choose some allocation which is strictly envy-free in net trades for every e^θ, while we have chosen a Walrasian allocation in each e^θ. Clearly, the latter is sufficient but not necessary. Moreover, since each agent does not know the realization of the other types in his cohort and θ, envy-freeness in every e^θ is itself not necessary for incentive compatibility. Thus, in general, there will be other allocations which satisfy all the criteria of our theorem. However, the allocation(s) constructed in the proof are special in that they are essentially the only ones which resemble separating rational expectations equilibria of e^θ. To see this, assume that each e^θ has a unique Walrasian equilibrium and that the Walrasian correspondence is continuous at this equilibrium with respect to perturbation of the density $P(\cdot|\theta)$. Thus, for r large and each realization sufficiently close to the expected realization of types given θ, every separating rational expectations equilibrium would yield allocations close to $x^\theta_{i_{t_s}r}$ for every type t_i of agent i. Hence, as r approaches infinity, the unique separating rational expectations equilibrium of e^r converges (in distribution) to x^r, the allocation in our proof.[11]

6. A special case of asymmetric information is one in which the information that any single agent possesses is redundant to the total information of the remaining agents regardless of the state of the world. In this case, it is known that any allocation can be a Bayes equilibrium outcome of some game (see, e.g., Blume and Easley (1983), Palfrey and Srivastava (1986), or Postlewaite and Schmeidler (1986)). In general, if not all information held by agents is redundant, there may not be any Pareto efficient allocation that is achievable as a Bayes equilibrium. Palfrey and Srivastava (1986) have investigated a similar

[10] In the proof, we choose some Walrasian allocation from each ε^θ. Thus if there is some ε^θ which has multiple Walrasian equilibria, our procedure will not yield a unique allocation for ε^r.

[11] In fact, if we could speak of a limit economy for $r = \infty$ such that the distribution of t's always matched the expected distribution given some θ, then the equilibria of our proof would correspond to the separating rational expectations equilibria, irrespective or the multiplicity of continuity of the Walrasian equilibria of ε^θ.

question to that dealt with in this paper. They considered a stochastic replication procedure for an economy that can be described roughly as follows. Consider a fixed finite set of states of the world and some set of partitions on that set of states. An agent's information is represented by a probability distribution over partitions of the set of states and a prior over the states themselves. For a given realization, an agent will know his partition of the states and the event in that partition in which the state of the world lies. The stochastic replication of an economy is then a sequence of independent draws of particular agents. Palfrey and Srivastava show that any allocation that is achievable with complete information can be asymptotically achieved when the number of replications goes to infinity.

The primary difference between our work and that of Palfrey and Srivastava is that they take the states of the world to be exogenously given. Since there are a finite number of partitions (by assumption), as they replicate, the probability that for the actual realization, any agent's information is redundant in the sense mentioned above goes to one. Palfrey and Srivastava use this to prove their result.

For the replication process in our model, as the number of agents increases, the number of states increases as well. Thus the possibility of achieving nearly efficient allocations comes about not because agents asymptotically have no private information, but because the private information they have becomes asymptotically unimportant in terms of the aggregate variables of the economy.

Mas-Colell and Vives (1991) develop and analyze a model that is much more similar to this paper. They consider, as we do, economies in which larger numbers of agents do not diminish the monopoly an individual agent has on some information. Mas-Colell and Vives construct a mechanism for private values economies (economies in which an agent's type enters only his own utility function) that yields the Walrasian equilibrium allocation for a continuum economy and prove that the Bayesian equilibrium allocation correspondence is upper hemicontinuous. All equilibria for sufficiently large economies must then be nearly efficient since the limit allocation is.

Their results differ from ours in that their approach allows one to say something about *all* Bayesian equilibrium allocations, whereas our results allow one to draw conclusions only about a particular allocation. On the other hand, their results deal with the private values case with no aggregate uncertainty, a more restrictive set of environments than dealt with in this paper.

7. Stiglitz and Greenwald (1988) show that in certain parameterized models, the outcome is inefficient for nearly all values of the parameters. This is, of course, not inconsistent with our results. The models that we examine have asymmetric information that generally precludes the possible attainment of efficient outcomes. Our results show that in some circumstances, this inefficiency asymptotically vanishes.

8. In the application of our model to the Akerlof lemons model, we pointed out that the natural way of replicating the lemons model was inconsistent with the assumptions of our theorem since in the natural replication process, the

number of commodities is increasing and the types of agents in one cohort enter the utility function of agents in other cohorts. One might suspect that it is the latter—the fact that an agent might have information that is relevant to all other agents—that was important, since the information of such an agent might be expected to have a significant impact on the price of a good even when there are many other agents. More specifically, one might conjecture that the asymptotic existence of an incentive compatible, individually rational, nearly efficient outcome might arise by replicating an economy in which agents' types did not enter other agents' utility functions. There is a difficulty with such a conjecture, however, as illustrated by the following example.

Consider a standard two sided asymmetric information problem (as in Myerson-Satterthwaite (1983)) which satisfies the condition that agents' types enter only their own utility function. In the two person case, there may be impossibility of designing incentive compatible mechanisms which guarantee outcomes that are both ex post individually rational and ex post efficient. If we replicate the problem by generating many pairs of agents with independent valuations (across pairs) and in which only a specific pair cares about the object they are to allocate, the inefficiency resulting from the asymmetric information may not disappear. In such a situation, the presence of additional pairs doesn't affect the problem facing any particular pair of allocating the good between the two agents. Of course, we have the number of commodities equal to the number of pairs here. If we were to put a uniform bound on the number of commodities for all numbers of agents, we would not be able to carry out the procedure. This example, in a sense, shows that restricting agents' types to enter only their own utility functions but still allowing the number of commodities to go to infinity will not assure that the inefficiency due to the asymmetry of information asymptotically vanishes.

The problem is that the replication leaves unchanged the importance of each agent's private information about the "value" of the relevant good. Fixing the number of commodities is a simple way of guaranteeing that each agent's demand for any good becomes small asymptotically. Presumably one could find alternative conditions that would guarantee this.[12]

9. A last comment is on the distinction between common value problems and private value problems. In our model private value problems are those in which agents had private information that was of direct relevance only to themselves. This would be represented by agents whose utility functions depended only on their own type. Common value problems could be represented in two ways. First, an agent's type might enter all other agents' utility functions. Both the lemons problem and the insurance problem had this feature. Our replication process excluded such problems.

There is a second way in which there might be common values, however. Agents' utility functions might depend on θ and types might be correlated to θ. In this way, agents' types may be of interest to other agents not because they

[12] We thank Martin Hellwig for helpful discussions on this issue.

enter directly into utility functions, but because they are correlated to θ which does enter. Standard finance problems are typically of this sort. Agents' information about a risky asset is of interest only in predicting unknown characteristics of the risky asset.

Note that this second way in which we might represent common values is consistent with our replication process. We would represent the set of unknown characteristics by Θ and replicate the set of agents by adding cohorts that have types (observe signals) that are correlated to θ. Thus, our theorem covers some problems in which there are common values. We should point out that we must make a regularity assumption that at least some agent's demand must depend on his type; this rules out the "pure" common value case in which no agents' types enter utility functions, only θ enters the utility functions. We consider this a topic for future work.

Graduate School of Business, Stanford University, Stanford, CA 94305, U.S.A.
and
Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104, U.S.A.

Manuscript received April, 1989; final revision received April, 1992.

APPENDIX

PROOF OF THEOREM: Let $e^\theta = \{\bar{u}_m, \bar{w}_m, C_m \text{ for } m \in M\}$ be an Arrow-Debreu economy where $M = N \times T$, $\bar{w}_{it} = w_t$ is the endowment of agent $(i,t) \in M$, $C_m = \boldsymbol{R}^l_+$ is the consumption set of m and $\bar{u}_m(x) = u_i(x, t, \theta)$ is the utility function of $m = (i,t)$. Furthermore, let there be a mass $P(t \mid \theta)$ of each agent (i,t) in the economy e^θ. Let $D^\theta_m(p)$ denote the demand of agent m in the economy e^θ. Note that the aggregate excess demand function z defined by

$$z(p) = \sum_m P(t \mid \theta)\left(D^\theta_m(p) - w_t\right)$$

satisfies all the conditions of Theorem 2, Chapter 2, p. 28 of Arrow and Hahn (1971). Thus there exists p^θ, x^θ, a Walrasian equilibrium of e^θ. That is, $x^\theta = (x^\theta_m)_{m \in M}$, $p^\theta \in \operatorname{Int} S^{l-1}$ such that

(1) $$E\left(z(p^\theta)\right) = 0;$$

(2) $$\bar{u}_m(x_m) > \bar{u}_m\left(x^\theta_m\right) \quad \text{implies} \quad p(x_m - \bar{w}_m) > 0.$$

We will consider two cases:

(a) For all $i \in N$, $t_i, t_i' \in T_i$ there exists t_{-i} and $\theta \in \Theta$ such that $x^\theta_m \neq x^\theta_{m'}$ for $m = (i, t_i, t_{-i})$ and $m' = (i, t_i', t_{-i})$.

(b) Condition (a) above fails.

We will prove the theorem for case (a) and then modify the proof by using the regularity assumption on p to prove case (b). This will be the only occasion in which the regularity assumption is used. Thus, if (a) is satisfied, then the regularity requirement is not needed.

First, for every $\lambda \in (0,1)$ define $(y^\theta_m)_{m \in M}$ as follows:

$$y^\theta_m = \begin{cases} \lambda x^\theta_m & \text{if} \quad x^\theta_m \neq \bar{w}_m; \\ \bar{w}_m & \text{otherwise.} \end{cases}$$

Given the strict concavity and monotonicity of u_i, it is clear that for λ sufficiently large but < 1:

(I) Every agent m either strictly prefers her allocation y^θ_m in e^θ to her endowment $\bar{w}_m$ or $y^\theta_m = \bar{w}_m$;

(II) either $\Sigma_m P(t \mid \theta)(y^\theta_m - \bar{w}_m) < 0$ in every coordinate or $y^\theta_m = \bar{w}_m$ for all m (this follows from the fact that $\bar{w}_m \in R^l_{++}$ for all m);

(III) either $m = (i,t)$ strictly prefers y^θ_m to $y^\theta_{m'}$ for $m' = (i, t')$ or $y^\theta_m = y^\theta_{m'}$.

Define the function F by

$$F(\beta, x) = \sum_m \beta_t (x_m - \bar{w}_m)$$

where $\beta = (\beta_t)_{t \in T}$ and $x = (x_m)_{m \in M}$.

Let $\beta^\theta = (P(t|\theta))_{t \in T}$. Thus (II) is equivalent to

(3) $$F(\beta^\theta, y^\theta) \ll 0 \quad \text{or} \quad y_m^\theta = \bar{w}_m \quad \text{for all } m.$$

By the continuity of F it follows that there exists $\delta > 0$ such that

(4) $$F(\beta, y^\theta) \leqslant 0 \quad \text{for all } \beta \text{ satisfying } \quad \|\beta^\theta - \beta\| \leqslant \delta.$$

Let $L(\beta) = \text{Min}_x\, p^\theta \cdot F(\beta, x)$ subject to $x \in R_+^l$ and $u_m(x) \geqslant \bar{u}_m(x_m^\theta) + \varepsilon/4$. Since x^θ is a Walrasian equilibrium, $L(\beta^\theta) > 0$. Furthermore, since L is continuous, there exists $\delta > 0$ such that

(5) $$L(\beta) > 0 \quad \text{for all } \beta \text{ such that } \quad \|\beta - \beta^\theta\| \leqslant \delta.$$

Define for all $r \geqslant 1, t \in T$, and $t^r \in T^r, f(t^r, t)$, the frequency of t in t^r by $f(t', t) = (1/r)\#\{s \leqslant r | t_s^r = t\}$ and $f(t^r) = f(t^r, t)_{t \in T}$ where $\#A$ denotes the cardinality of the set A.

Let $T_{r\delta}^*(\theta) = \{t^r \in T^r | \|f(t^r) - \beta^\theta\| < \delta\}$ and $T_{r\delta}^* = U_\theta T_{r\delta}^*(\theta)$. By the nontriviality (iii) of P for all $\delta > 0$ small enough,

(6) $$T_{r\delta}^*(\theta) \cap T_{r\delta}^*(\theta') \neq \phi \quad \text{implies} \quad \theta = \theta'.$$

Now we are ready to define the desired allocation for the random replica economy e^r:

$$x_{ts}^r(t^r) = \begin{cases} w_t & \text{if } t^r \notin T_{r\delta}^*; \\ y_{t_s^r}^\theta & \text{if } t^r \in T_{r\delta}^*(\theta). \end{cases}$$

To show (F), (IC), (IR), and (X_εER), we will argue that for each of these conditions one can choose $\bar{\delta}$ appropriately such that, for $\delta \in (0, \bar{\delta})$, choosing r large enough establishes the desired conclusion.

Feasibility: For any given λ choose δ such that (4) is satisfied. If $t^r \notin T_{r\delta}^*$ then every agent receives her own endowment; hence, there is nothing to prove. If $t^r \in T_{r\delta}^*(\theta)$ then

$$\sum_t \sum_s (x_{ts}^r(t^r) - w_t) = r \sum_t \sum_t f(t^r, t)(y_{tt}^\theta - w_t)$$
$$= r \cdot F(f(t^r), y^\theta) \leqslant 0$$

by (4).

Incentive Compatibility: Note that

(7) $$u_{ts}(x^r | t_{ts}) = \sum_\theta P(\tilde{t}^r \in T_{r\delta}^*(\theta) | t_{ts}) \cdot \sum_{\theta'} P(\theta' | t^r \in T_{r\delta}^*(\theta) \quad \text{and } t_{ts})$$
$$\cdot \sum_{t_{-t}} P(\tilde{t}_s^r = t | \tilde{t}^r \in T_{r\delta}^*(\theta), \theta', t_{ts}) \cdot u_t(y_{tt}^\theta, t, \theta') + [1 - P(t^r \notin T_{r\delta}^*)] \cdot A$$

where $t = (t_{-t}, t_{ts})$, and A is some weighted average of the utilities of w_t in various states (t, θ). (Note that A is bounded over all r.) Thus, according to x^r, there are two sources of uncertainty. First, there is the probability that the observed frequency of t's will be close to the expected frequency given some $\theta[P(\tilde{t}^r \in T^r(\theta) | t_{ts})]$. Second, there is the probability that the θ identified in this manner will (or will not) be the correct one. It follows from the law of large numbers that as r approaches infinity $P(t^r \in T_{r\delta}^*(\theta) | \theta$ and $t_{ts})$ approaches 1. Therefore, $P(\tilde{t}^r \in T_{r\delta}^*(\theta) | t_{ts}) = P(\tilde{t}^r \in T_{r\delta}^*(\theta) | \theta$ and $t_{ts}) \cdot P(\theta | t_{ts})$ approaches $P(\theta | t_{ts})$. Furthermore, by Bayes' law

$$P(\theta' | t^r \in T_{r\delta}^*(\theta) \text{ and } t_{ts}) = \frac{P(\tilde{t}^r \in T_{r\delta}^*(\theta) | \theta' \text{ and } t_{ts}) \cdot P(\theta' | t_{ts})}{P(\tilde{t}^r \in T_{r\delta}^*(\theta) | t_{ts})}.$$

Thus $P(\theta'|t^r \in T^*_{r\delta}(\theta)$ and $t_{is})$ approaches 0 as $r \to \infty$ for $\theta' \neq \theta$ (by (6)). Therefore, (7) becomes

$$u_i(t^r|t_{is}) = \sum_{\theta} P(\theta|t_i) \sum_{t_{-i}} P(\tilde{t}_{-i}|\theta, t_{is}) u_i(y^{\theta}_{it}, t, \theta)$$

where $t = (t_{-i}, t_{is})$. Note that this is a weighted average of $u_i(y^{\theta}_{it}, t, \theta) = \bar{u}_{it}(y^{\theta}_{it})$. By hypothesis (a), agent t_{is} will, by reporting $t'_i \neq t_i$, affect his allocation in at least one state (t_{-i}, θ). Furthermore, by (I) above, the agent will receive a less desirable outcome and since $P(\theta|t_i)$ and $P(t_{-i}|\theta, t_{is})$ are all nonzero (by the full support assumption on P), this will lead to a strictly lower utility for misrepresentation in the limit and hence also for r large enough. Since there are a finite number of players and types for each player, we can choose r large enough that this is true for all i and t_i. This establishes incentive compatibility.

(*Ex-post*) *Individual Rationality*: The ex-post utility of x^r for any agent i, t_{is} is given by

$$\sum_{\theta} P(\theta|t^r) u_i(x^r(t^r), t_s, \theta). \tag{8}$$

If $t^r \notin T^*_{r\delta}$, then clearly x^r is IR since $x^r(t^r) = w_i$. Suppose $t^r \in T^*_{r\delta}(\theta)$ for some θ. Again, if $x^r(t^r) = w_i$, there is nothing to prove. If not, then by (III) we know that $u_i(x^r(t^r), t_s, \theta) > u_i(w_i, t_s, \theta)$. Hence, the proof hinges on showing that $P(\theta|t^r)$ can be made arbitrarily close to 1 for each $t^r \in T^*_{r\delta}$ (note that this is a stronger statement than the one made in proving IC above where it was shown that by the law of large numbers $P(\theta|t^r \in T^*_{r\delta}(\theta))$ approaches 1).

We will show that for all sufficiently small $\varepsilon' > 0$ there exists $\bar{\delta}$ and $\bar{r}_{\bar{\delta}}$ such that $t^r \in T^*_{r\delta}(\theta)$ for $\delta \in (\bar{\delta}, 1)$ and $r \geqslant \bar{r}_{\bar{\delta}}$ implies $P(\theta|t^r) > 1 - \varepsilon'$. This, together with the boundedness of u_i, will establish the desired conclusion.

By Bayes' Law

$$P(\theta|t^r) = \frac{P(\theta)P(t^r|\theta)}{P(t^r)} = \frac{P(\theta)P(t^r|\theta)}{\sum_{\hat{\theta} \in \Theta} P(\hat{\theta})P(t^r|\hat{\theta})}.$$

By the multinomial formula

$$P(t^r|\hat{\theta}) = \frac{r!}{\prod_{t \in T}(rf(t^r, t))!} \prod_{t \in T} P(t|\hat{\theta})^{rf(t^r, t)}.$$

Let $H(\beta, \beta') = \prod_{t \in T} \beta_t^{\beta'_t}$ for $\beta = (\beta_t)_{t \in T}$ and $\beta' = (\beta'_t)_{t \in T}$. Hence

$$P(t^r|\theta) = \frac{r!}{\prod_{t \in T}(rf(t^r, t))!} H(\beta^{\theta}, f(t^r))^r$$

and

$$1/P(\theta|t^r) = 1 + \sum_{\hat{\theta} \in \Theta, \hat{\theta} \neq \theta} \frac{P(\hat{\theta})}{P(\theta)} \cdot \left(\frac{H(\beta^{\hat{\theta}}, f(t^r))}{H(\beta^{\theta}, f(t^r))} \right)^r. \tag{9}$$

Observing that $H(\cdot, f(t^r))$: $S^{|T|-1} \to R$ is continuous and has a unique maximum at $\beta = f(t^r)$ establishes that for δ sufficiently small and $t^r \in T^*_{r\delta}(\theta)$, $H(\beta^{\theta}, f(t^r)) > H(\beta^{\hat{\theta}}, f(t^r))$ for $\hat{\theta} \neq \theta$. Thus for r large enough, all the terms after the summation in (9) above are close to zero; hence, $P(\theta|t^r)$ is close to 1 as desired.

(*Ex-post*) *ε-efficiency*: Consider v^r such that for all i, s

$$\begin{aligned} \sum_{\theta} u_i(v^r_{is}(t^r), t^r_s, \theta) \cdot P(\theta|t^r) &> \sum_{\theta} u_i(x^r_{is}(t^r), t^r_s, \theta) P(\theta|t^r) + \frac{\varepsilon}{4} \\ &= \sum_{\theta} u_i\left(y^{\theta}_{i t^r_s}, t^r_s, \theta\right) P(\theta|t^r) + \frac{\varepsilon}{2}, \end{aligned} \tag{10}$$

for all $t^r \in T^*_{r\delta}(\theta)$ by (5).

Note that since $P(t^r \in T^*_{r\delta}) > 1 - \delta$ for $r > \bar{r}_{\delta}$ then for $\delta < \varepsilon$, $T^*_{r\delta}$ can serve as the set E in the definition. Since u_i's are bounded and the fact that $P(\theta|t^r) > 1 - \delta$ for $r > \bar{r}_{\delta}$ and $t^r \in T^*_{r\delta}(\theta)$

implies that for sufficiently small δ and some $\bar{r}_\delta$ then

$$(11) \qquad u_i(v_{is}^r(t^r), t_s^r, \theta) > u_i\left(y_{it_s^r}^\theta, t_s^r, \theta\right) + \varepsilon.$$

Define $v^* = (v_{it}^*)_{it \in N \times T}$ by $v_{it}^* = f(t^r, t)\Sigma_{ss.t.\ t_s^r = t} v_{is}^r(t^r)$.

Note that by the concavity of $u_i(\cdot, t, \theta)$, (11) implies that for δ sufficiently small and $r > \bar{r}_\delta$,

$$(12) \qquad u_i(v_{it}^*, t, \theta) > u_t\left(y_{it}^\theta, t, \theta\right) + \varepsilon.$$

Furthermore,

$$(13) \qquad \begin{aligned} \sum_i \sum_s (v_{is}^r(t^r) - w_i) &= \sum_t \sum_{ss.t.\ t_s^r = t} \sum_i (v_{is}^r(t^r) - w_i) \\ &= \sum_t \sum_{ss.t.\ t_s^r = t} \sum_i r \cdot f(t^r, t)(v_{it}^* - w_t) \\ &= r \cdot F(f(t^r), v^*). \end{aligned}$$

But $f(t^r) \in B_\delta(\beta^\theta)$ whenever $t^r \in T_{r\delta}^*(\theta)$ so by (5)

$$p^\theta \cdot F(f(t^r), v^*) > 0;$$

hence there exists $v \in \{1, 2, \dots l\}$ such that

$$F_v(f(t^r), v^*) > 0$$

which by (13) implies z^r is not feasible and completes the proof.

All that remains to be shown is that if the hypothesis (a) is not satisfied (i.e., in case (b)), then the regularity assumption can be used to modify the proof above in the appropriate manner. Notice that a problem arises only with the proof of incentive compatibility since the other proofs did not utilize hypothesis (a). Here the difficulty arises from the fact that if types t_i and t_i' ($t_i \neq t_i'$) of agent i always receive the same allocation for every one of the other agents, then in the limit (as $r \to \infty$), t_i (and t_i') will be exactly indifferent between announcing t_i and t_i'. But then for any $r < \infty$ the effect of this announcement on the event that the observed t^r will (or will not) belong to $t^r \in T_{r\delta}^*$ cannot be ignored. Choose θ such that $y_{it}^\theta \neq w_t$ for some t_{-i} and $t = (t_{-i}, t_i)$ (if there is no such θ, there is nothing to prove). Thus the first part of (3) is satisfied with strict inequality. By the regularity assumption, there exists a price $p' \neq p^\theta$ arbitrarily close to p^θ such that the demands of $m = (i, (t_i, t_{-i}))$ and $m' = (i, (t_i', t_{-i}))$ are distinct. Furthermore, by the continuity of the demands, this price can be chosen close enough to p^θ so that some of the demands still satisfy I–III above. Replacing y_m^θ with the demand of m at p' and $y_{m'}^\theta$ with the demand of m' at p' enables us to complete the proof as in the case of hypothesis (a). A similar argument works if there are more than two types who have exactly the same demand for every t_{-i} and θ.

REFERENCES

AKERLOF, G. (1970): "The Market for Lemons," *Quarterly Journal of Economics*, 89, 488–500.

ARROW, K., AND F. HAHN (1971): *General Competitive Analysis*. San Francisco, CA: Holden-Day. Now distributed by North Holland, Amsterdam.

AUMANN, R. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18.

BILLINGSLEY, P. (1963): *Convergence of Probability Measures*. New York: Wiley.

BLUME, L., AND D. EASLEY (1983): "On the Game-Theoretic Foundations of Market Equilibrium with Asymmetric Information," mimeo, Cornell University.

GRESIK, T., AND M. SATTERTHWAITE (1989): "The Rate at Which a Simple Market Converges to Efficiency as the Number of Traders Increases: An Asymptotic Result for Optimal Trading Mechanisms," *Journal of Economic Theory*, 48, 304–332.

MAS-COLELL, A., AND X. VIVES (1991): "Implementation with a Continuum of Agents," mimeo, Harvard University.

MYERSON, R., AND M. SATTERTHWAITE (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265–281.

PALFREY, T., AND S. SRIVASTAVA (1986): "Private Information in Large Economies," *Journal of Economic Theory*, 39, 34–58.

Postlewaite, A., and D. Schmeidler (1986): "Implementation in Differential Information Economies," *Journal of Economic Theory*, 39, 14–33.

Roberts, D. J., and A. Postlewaite (1976): "The Incentives for Price-Taking in Large Economies," *Econometrica*, 44, 115–128.

Satterthwaite, M., and S. Williams (1989): "The Rate of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large," *Review of Economic Studies*, 56, 477–498.

Stiglitz, J., and B. Greenwald (1988): "Externalities in Economies with Imperfect Information and Incomplete Markets," *Quarterly Journal of Economics*, 101, 229–264.

Wilson, R. (1987): "On Equilibria of Bid-Ask Markets," in *Arrow and the Ascent of Modern Economic Theory*, edited by George Feiwel. London: Macmillan Press.