JOURNAL OF ECONOMIC THEORY 52, 208–222 (1990)

Notes, Comments, and Letters to the Editor

Continuous Implementation of Constrained Rational Expectations Equilibria*

DAVID WETTSTEIN

University of Western Ontario,
London, Ontario, Canada N6A 5C2

Received December 22, 1987; revised February 14, 1990

We investigate the implementability of the Social Choice Correspondence induced by Constrained Rational Expectations Equilibria. We construct an "almost continuous" mechanism that implements it, provided that information held by each agent is non-exclusive *Journal of Economic Literature* Classification Number: 026.

1. INTRODUCTION

The theory of implementation deals with ways by which a society can achieve desired outcomes. A distinctive feature of the theory is its alertness to the possibility that individuals behave strategically when asked to provide some information or take a particular course of action.

This paper is on the theory of implementation in economic environments with incomplete information. The economic environment is fully described by preferences and initial endowments of the individuals constituting the economy. An outcome is some reallocation of the initial endowments. The outcomes deemed desirable by the society are represented by a Social

* The research reported in this paper constitutes part of my Ph.D. thesis submitted to the University of Minnesota. I wish to thank my thesis advisor Professor Leonid Hurwicz for helpful discussions. James Jordan provided several comments on an earlier version. Helpful conversations with Andrew Postlewaite and David Schmeidler are gratefully acknowledged. Long talks with Abhijit Sengupta have considerably improved the final product. I wish to thank an associate editor for numerous comments and suggestions that have removed several ambiguities and greatly improved the exposition and formulation of the results. Two anonymous referees made several helpful remarks. Debbie Fretz provided able editorial assistance.

0022-0531/90 $3.00

Choice Correspondence (SCC). (An SCC is a collection of allocation rules; an allocation rule maps environments into outcomes.)

A central designer is provided with the Social Choice Correspondence as well as information regarding the environment. His task is to design a mechanism implementing the SCC, i.e., a mechanism whose set of equilibrium outcomes for each environment coincides with the set of outcomes prescribed by the SCC for it. The SCC is called implementable if such a mechanism exists. It is perfectly possible that no such mechanism exists.

The problems arise owing to the fact the designer might not have all the relevant information and would have to rely on the individuals to supply it. However, realizing their messages determine the outcome reached, the individuals might find it advantageous to provide false information. Thus individuals need to be given the correct incentives to reveal any private information they might have.

The question of implementability hinges crucially on the equilibrium notion used, and we consider the Bayesian–Nashian Equilibrium (BNE) introduced by Harsanyi [3] for games with incomplete information. Postlewaite and Schmeidler [10, 11] provided a set of sufficient conditions for the implementability of a given SCC. In later work Palfrey and Srivastava [9] addressed the same issue under a less restrictive informational structure.[1] Both proofs constructed particular mechanisms carrying out the implementation.

Due to the general nature of the SCC's under consideration in those works, the implementing mechanisms turn out to be quite complex. A major drawback is the fact that the designer is required to know individual preferences. Furthermore these mechanisms are highly discontinuous and they assume that the information structure is discrete[2] (there are a finite number of states of the world).

Discontinuous mechanisms suffer from the fault that a slight mistake might cause a considerable change in the outcome reached. A continuous mechanism also has the advantage that it is robust with respect to some misspecifications. More on the importance of continuity can be found in Postlewaite and Wettstein [12] and Wettstein [14].

One of the leading models used to characterize the competitive behavior of agents under uncertainty is Rational Expectations Equilibria (REE).

[1] They did not assume information is non-exclusive. This Non-Exclusivity of Information (NEI) assumption states that the pooled information of any $n-1$ individuals, where n is the total number of individuals, is complete.

[2] It is true that the assumption, made by the above mentioned authors, regarding the discreteness of the information structure, renders continuity with regard to information states trivial. However, with respect to the other dimensions of the strategy space, those related to allocations and other physical parameters, these mechanisms are not continuous. It is natural to ask for continuity in the outcome function in all dimensions as we shall do.

REE are a close counterpart of the Walrasian Equilibria in economies with complete information. A good discussion of REE with relevant references can be found in Jordan and Radner [7]. With a finite number of agents the REE correspondence is not immune to strategic behavior on the part of the agents in the economy. A natural question is whether there exist institutions circumventing this problem or in other words: Is the REE correspondence implementable? Without further restrictions the answer is no.[3]

Palfrey and Srivastava [8] have demonstrated that the REE correspondence is implementable when restricted to a subset of economies where all equilibrium allocations are interior. Their approach is to appeal to the sufficient conditions which guarantee implementability in the case of a general correspondence. Thus their approach shares the same drawbacks mentioned with regard to the general mechanism: requirement of knowledge of individual preferences and discontinuity. This leaves open the question whether this particular correspondence can be implemented via a simpler mechanism with nicer properties.

The same question has been raised regarding the Walrasian correspondence in economies with complete information and was answered in part by Hurwicz [5], Schmeidler [13], and Postlewaite and Wettstein [12].

Rather than restricting the set of economies, we shall show that a modified version of the REE, namely constrained Rational Expectations Equilibria (CREE), where an individual is constrained to demand a bundle that does not exceed aggregate supply, is indeed implementable under certain conditions. Our main result is that there exists a feasible and almost continuous[4] mechanism, which under certain conditions implements the CREE. Our mechanism, in contrast to those previously suggested, does not require the designer to know the individual preferences. The information structure would be continuous and the strategy space itself will be reminiscent of a market in which prices and quantities are announced. The mechanism provides a feasible outcome for any configuration of actions chosen by the individuals (both in and out of equilibrium).

The timing in our mechanism is somewhat unconventional. We assume the individuals start out with symmetric information and can take an action before they observe their private information. This initial stage can be dispensed with at the cost of technical measurability assumptions. This will be discussed further in connection with the construction of our

[3] An example for non-implementability of the Walrasian correspondence can be found in Hurwicz, Maskin, and Postlewaite [6]. Similar examples can be worked out for the REE correspondence.

[4] By almost continuous we mean that all the operations carried out by the outcome function, with the possible exception of some projections (see footnote 6), are continuous.

mechanism. After individuals observe their private information, there is a signalling stage. This stage is natural in modelling market behavior under conditions of incomplete information and greatly simplifies the presentation.

The most important and restrictive of our assumptions is that agents do not have exclusive information. The Non-Exclusive Information (NEI) assumption appeared in Postlewaite and Schmeidler [10, 11]. In related work Blume and Easley [2] have shown that violation of the NEI leads to non-implementability of the REE.

The outline of the paper is as follows: The second section describes the environments we consider and introduces notation and definitions. In the third section we construct a feasible and almost continuous mechanism implementing the CREE. The fourth and final section discusses possible extensions and further lines of research.

2. Notation and Definitions

(R^n, B^n, μ)—The n-dimensional Euclidean space viewed as a probability space with B^n denoting the set of Borel sets and μ a probability measure defined on B^n. A generic element of this space will be denoted by $\varepsilon = (\varepsilon_1, ..., \varepsilon_n)$, where ε_i is in R for all $i = 1, ..., n$.

J—The support[5] of μ. It is interpreted as the set of possible states of the world. J is assumed to be compact.

$J_{i_1 \cdots i_m}$—The projection[6] of J onto the $(i_1, ..., i_m)$ axis. Since J is compact this would be a closed set as well.

The economy E has n individuals and k commodities. It is specified by the following:

1. An n-tuple of Von Neumann–Morgenstern utility functions $u_i: R^k_+ \times J \to R$, where $u_i(x, \varepsilon)$ denotes the satisfaction individual i derives from consuming commodity bundle x, when the state of the world is ε. The u_i's are assumed to be measurable[7] in their last n coordinates.

2. An n-tuple of initial endowments $(w_1, ..., w_n)$, where w_i is in R^k_+ for all $i = 1, ..., n$.

The aggregate endowment is $W = \sum_{i=1}^n w_i$.

[5] The support of μ is the smallest closed set in R^n with measure 1.

[6] Given a point x and a set S, the projection of x on the set S is the point y in S closest to x.

[7] Unless otherwise specified, measurability is taken to be in the Borel sense.

Note that initial endowments do not vary with the state of the world. If one were to consider non-feasible implementation (i.e. the use of mechanisms that might, for certain configurations of actions by individuals, yield non-feasible outcomes), it would have been enough to assume the endowments were common knowledge among the individuals as in Aumann [1]. The designer would not need to know endowments as he would only be looking at net trade rules. To show this we can proceed as in Hurwicz, Maskin, and Postlewaite [6], who considered the issue under complete information.

We assume that the structure of the economy and the probability space are common knowledge. Individual i observes ε_i privately, which leads to asymmetric information. When one considers a finite set of states of the world this way of modelling the uncertainty is equivalent to the usual formulation using partitions over a given set of states. The number of distinct values for ε_1 will equal the number of elements in the partition of individual 1 and so forth. Our way of formulating the asymmetric information is better suited for studying a continuous information structure.

A—The set of feasible allocations for E:

$$A = \left\{ (x_1, ..., x_n) \in R_+^{nk} \;\middle|\; \sum_{i=1}^{n} x_i^j \leqslant \sum_{i=1}^{n} w_i^j \text{ for } j = 1, ..., k \right\}.$$

f—An allocation rule defined as a measurable function associating a feasible allocation with each ε. $f = (f_1, ..., f_n)$ where f_i denotes the bundle allocated to individual i.

F—A Social Choice Correspondence (SCC) defined as a collection of allocation rules.

The way we define an SCC deserves some comment. One could define an SCC as a mapping associating a set of good allocations (not allocation rules) with each state of the world. Our definition, which has been commonly used in this literature, is more specialized. An allocation rule reflects the idea that there is a certain property we want an allocation of goods to satisfy for all states of the world. This definition is motivated by the timing of uncertainty in our model. The individual has to take an action observing his private information. This information is usually not sufficient to determine the precise state of the world, and hence the action the individual takes may have different consequences for different states, which he cannot distinguish at the time he takes the action. Our allocation rule incorporates this measurability constraint.

Under conditions of incomplete information, one must allow the possibility that the prices at which trade takes place partially (or completely) reveal the state of the world. This notion is captured in the following definition of a Rational Expectations Equilibrium.

A pricing function $P: J \to R^k_+$ and an allocation rule f constitute a Rational Expectations Equilibrium (REE) if the following conditions hold:

(i) P is measurable.

(ii) f_i is measurable with respect to ε_i and $P(\varepsilon)$[8] for all $i = 1, ..., n$.

(iii) For all $i = 1, ..., n$ and almost everywhere (a.e.) in J, $f_i(\varepsilon)$ solves

$$\begin{aligned} &\operatorname*{Max}_{x_i} E(u_i(x_i, \varepsilon) \mid \varepsilon_i, P(\varepsilon)) \\ &\text{S.T.} \\ &\qquad P(\varepsilon) \cdot x_i \leqslant P(\varepsilon) \cdot w_i \\ &\qquad x_i \geqslant 0 \end{aligned}$$

(iv) $P(\varepsilon) \cdot (W - \sum_{i=1}^n f_i(\varepsilon)) = 0$ a.e. in J.

The expectations are conditional on the probability measure on J given that the ith coordinate equals ε_i and P equals $P(\varepsilon)$.

As noted earlier the REE correspondence is not implementable. This is taken care of by limiting ourselves to a Constrained Rational Expectations Equilibrium. $P: J \to R^k_+$ and an allocation rule f constitute a CREE if the following holds:

(i), (ii), and (iv) as in the definition of an REE and (iii) replaced by:

(iii)′ For all $i = 1, ..., n$ and a.e., in J, $f_i(\varepsilon)$ solves

$$\begin{aligned} &\operatorname*{Max}_{x_i} E(u_i(x_i, \varepsilon) \mid \varepsilon_i, P(\varepsilon)) \\ &\text{S.T.} \\ &\qquad P(\varepsilon) \cdot x_i \leqslant P(\varepsilon) \cdot w_i \\ &\qquad x_i \leqslant W \\ &\qquad x_i \geqslant 0 \end{aligned}$$

Thus the difference between CREE and REE is that in the former individuals are not allowed to demand a commodity bundle which exceeds in one or more of its coordinates the aggregate endowment vector. It is obvious that both definitions coincide for interior solutions.

Before outlining the specific mechanism constructed in this paper, we formally define the notions of mechanism, Bayesian–Nashian Equilibrium (BNE), and implementation.

[8] Thus the values f_i assumes for any ε in J can depend only on the ith coordinate of ε and the associated price vector.

A mechanism G consists of:

1. An n-tuple of measurable strategy sets $(S_1, ..., S_n)$,

$$S_i = B_i \times H_i \times D_i,$$

where B_i denotes acts that have to be taken before the observation of any private information, H_i denotes acts taken after the observation of the private information, but prior to observing any signals, and D_i denotes acts taken after observing a message sent by the designer. We write

$$S = \prod_{i=1}^{n} S_i; \qquad B = \prod_{i=1}^{n} B_i; \qquad H = \prod_{i=1}^{n} H_i; \qquad D = \prod_{i=1}^{n} D_i$$

2. A measurable signalling structure $P = (P_1, ..., P_n)$, where $P_i: B \times H \to \Gamma_i$. P_i is the signal received by the ith individual and Γ_i is a measurable signal space,

$$\Gamma = \prod_{i=1}^{n} \Gamma_i.$$

3. A measurable outcome function

$$g: B \times H \times D \to A,$$

$g = (g_1, ..., g_n)$ where g_i denotes the bundle received by individual i.

A strategy for individual i consists of a choice b_i in B_i, a measurable function $h_i: J_i \to H_i$, and a measurable function $d_i: J_i \times \Gamma_i \to D_i$.

In defining a Bayesian–Nashian Equilibrium of such a mechanism we use the following notation: if $s = (s_1, ..., s_n)$, then for $\hat{s}$ in S_i we write $(s_{-i}, \hat{s}) = (s_1, ..., s_{i-1}, \hat{s}, s_{i+1}, ..., s_n)$. A BNE of the mechanism is an n-tuple of strategies

$$\tilde{s} = (\tilde{s}_1, ..., \tilde{s}_n), \qquad \text{where} \quad \tilde{s}_i = (\tilde{b}_i, \tilde{h}_i, \tilde{d}_i) \qquad \text{satisfies for all} \quad i = 1, ..., n$$

$$E(u_i(g_i(\tilde{s}), \varepsilon)) \geq E(u_i(g_i(\tilde{s}_{-i}, \hat{s}), \varepsilon)) \qquad \text{for all } \hat{s} \text{ in } S_i.$$

Note that those imply that the strategy choice made for any ε_i and any signal subsequently observed is optimal given the choices made by others, except possibly on a set of ε's and signals of measure zero.

The strategies chosen by individuals yield an allocation rule a; $a(\varepsilon) = g(s)$.

Let $N_G(E)$ be the set of allocation rules corresponding to BNE of the mechanism for E.

We say that G implements F if

$$N_G(E) = F \qquad \text{a.e.}^{9}$$

The use of an explicit signalling stage is not essential; the same mechanism can be represented by a more complicated one which involves no signalling. This can be done in the same way in which one moves from a game in extensive form to a game in strategic form. However, the use of a signalling stage is quite natural in the context of an REE.

As noted earlier, to achieve implementability we must assume that the support of μ satisfies a Non-Exclusive Information (NEI) assumption. It can be formulated in various ways. We formulate it as follows:

(NEI) $\sum_{i=1}^{n} \varepsilon_i = 0$ for all ε in J.

Admittedly this is not the most general formulation possible; however, all that is subsequently used is the fact that any $(n-1)$ coordinates of ε uniquely determine the remaining one.

3. The Mechanism

The Strategy Spaces

The strategy space for individual i is defined to be $S_i = B_i \times H_i \times D_i$, where $B_i = \tilde{P}$, the set of all measurable functions from J into R^k_+; $H_i = J_i$; $D_i = R^k \times R_{++}$; with generic element (P^i, r_i, z_i, m_i); $P_i \in B_i$, $r_i \in H_i$, $(z_i, m_i) \in D_i$.

The informational constraints imply that P^i is chosen before any private information is observed, r_i picked after the observation of ε_i and (z_i, m_i) decided upon after observing a message, which in our case will be some price vector.

Individual i's choices can be intuitively interpreted as a suggested pricing rule (P^i), the ε_i he "observed" (r_i), a net trade requested (z_i), and finally a parameter (m_i) used in the process of averaging the possibly conflicting demands of all the individuals.

[9] Any allocation rule in $N_G(E)$ is equal to some allocation rule in F almost everywhere, and vice versa.

The Outcome Function

Step 1. We construct a weighted average of the P^i's. To this end, first define, for ε in J,

$$\alpha_j(\varepsilon) = \sum_{t,t' \neq j} |P^t(\varepsilon) - P^{t'}(\varepsilon)|,$$

where P^t denotes the pricing function sent by individual t.

Now define

$$\alpha = \sum_{i=1}^{n} \alpha_j$$

$$\beta_j = \begin{cases} \dfrac{\alpha_j}{\alpha} & \text{if} \quad \alpha > 0 \\ \dfrac{1}{n} & \text{if} \quad \alpha = 0 \end{cases} \qquad j = 1, ..., n;$$

finally, let

$$\bar{P}(\varepsilon) = \sum_{j=1}^{n} \beta_j P^j(\varepsilon).^{10}$$

It is clear from the construction of $\bar{P}$ that if all the individuals announce the same pricing function then no single individual can change $\bar{P}$ by changing his own announcement.

One could assume the individuals send in their pricing functions after having observed their private information. This would eliminate the need for the existence of a stage with symmetric information. In principle if each individual knew the strategies (part of which would be a mapping from their observation into the set of measurable pricing functions) of all the other individuals, he would be able to compute the pricing function the mechanism would construct for each possible state of the world. This collection of functions would form the pricing function (the same for all the individuals) which would be used to form conditional expectations, given some announced price vector. In equilibrium this would turn out to be part of an REE, and the results reported here would continue to hold. This additional stage, however, complicates the presentation and necessitates the assumption that the pricing function calculated in that way is measurable. To simplify the exposition, we assume the individuals start out with symmetric information, at which point they are asked to send in pricing functions.

[10] We note that $\bar{P}(\varepsilon)$ is a continuous function of the P_j's even though the β_j's are not.

Step 2. This stage will involve several projections and care must be taken to ensure all these operations are measurable. We need to project tuples of numbers announced by the individuals on n closed sets in R^{n-1}. Prior to starting the operation of the mechanism, n measurable functions (assumed to be common knowledge among the individuals) carrying out these projections are constructed. When the sets in question are convex the projection operation yields a continuous and, a fortiori, a measurable function. If the sets are not convex the projection operation turns out to be a correspondence; however, by Lemma 1 in Hildenbrand [4, p. 55], the correspondence admits a measurable selection.

The functions are used by the mechanism to construct n profiles in J. The ith profile denoted by $(\tilde{r}_1^i, ..., \tilde{r}_n^i)$ is constructed in the following manner:

The projection (by the above mentioned function) of $(r_1, ..., r_{i-1}, r_{i+1}, ..., r_n)$ on $J_{1,...,i-1,i+1,...,n}$ is denoted by $(\tilde{r}_1^i, ..., \tilde{r}_{i-1}^i, \tilde{r}_{i+1}^i, ..., \tilde{r}_n^i)$ and $\tilde{r}_i^i$ is defined by $\tilde{r}_i^i = -\sum_{j \neq i} \tilde{r}_j^i$. The profile $\tilde{r}^i = (\tilde{r}_1^i, ..., \tilde{r}_n^i)$ is indeed in J. The following two facts are evident from the construction:

(i) If $(r_1, ..., r_n)$ is in J then all the n-profiles are identical and equal to r.

(ii) A change in the strategy of individual i will not change the ith profile.

By now we have a pricing function $\bar{P}$ and n profiles $(\tilde{r}^i)_{i=1}^n$. At this point individual i is told what $\bar{P}(\tilde{r}^i)$ is, i.e., he is told a certain price vector in R_+^k.

Step 3. Informed of the relevant price vector, each individual announces (z_i, m_i).

Let

$$K^i(\bar{P}(\tilde{r}^i), w_i, W) = \{z \in R^k \mid \bar{P}(\tilde{r}^i) \cdot z = 0, z + w_i \geqslant 0, z + w_i \leqslant W\}$$

and let y_i be the closest point in K^i to z_i. Note that K^i is non-empty, closed, and convex. Define

$$C = \left\{ m \in R_{++} \mid m \cdot m_i \leqslant 1 \text{ for all } i = 1, ..., n \text{ and } m \cdot \sum_{i=1}^{n} m_i(y_i + w_i) \leqslant W \right\}$$

and let $m^* = \max_{m \in C} m$.

Finally a projection on J is called for. As before we use a measurable projection function, constructed prior to the game and assumed to be common knowledge. Project $(r_1, ..., r_n)$ on J, get a point ε^* in J, and let

$$g_i(s_1, ..., s_n) = m^* \cdot m_i \cdot t_i(y_i + w_i),$$

where

$$t_i = \left(1 + \frac{\sum_{j=1}^{n} m_j}{m_i} |r - \varepsilon^*|\right)^{-1}.$$

THEOREM. *Given the assumptions*

(A1) $n \geqslant 3$

(A2) *for all* $i = 1, ..., n$, u_i *is continuous, strictly increasing, and strictly concave in its first* k *arguments for any* ε *in* J

(A3) *for all* $i = 1, ..., n$, $u_i(W, \varepsilon)$ *is integrable in its last* n *arguments*

(A4) *for all* $i = 1, ..., n$, w_i *is in* R^k_{++}

(A5) J *satisfies* NEI

the mechanism constructed above implements the CREE.

Proof. In the first part of the proof we show that for any CREE allocation rule there exists a BNE allocation of this mechanism that coincides with it a.e.

Let $P(\varepsilon)$ and $f(\varepsilon)$ denote a CREE. We now construct the following BNE which yields f:

$P^i = P$		for all $i = 1, ..., n$
$r_i = \varepsilon_i$		for all ε in J and all $i = 1, ..., n$
$z_i = f_i(\varepsilon) - w_i$;	$m_i = 1$	for all ε in J, all possible signals arising out of the choice of P and r, and all $i = 1, ..., n$.

The first step yields P and no single individual can change it by deviating and declaring another pricing function.

The second step yields the true profile $(\varepsilon_1, ..., \varepsilon_n)$ for all the individuals, and each individual is told the true price

$$P(\varepsilon_1, ..., \varepsilon_n).$$

For a.e. ε in J, z_i belongs to K^i and hence $y_i = z_i$ a.e. in J. Since f was a CREE we have that for a.e. ε in J: $\sum_{i=1}^{n} f_i(\varepsilon) = W$, and hence for a.e. ε in J, $\sum_{i=1}^{n} (y_i + w_i) = W$. Since $m_i = 1$ for all $i = 1, ..., n$, we get $m^* = 1$ for a.e. ε in J. Further, since individuals report their true ε_i, $t_i = 1$ for all ε in J. Thus $g(s_1, ..., s_n) = f(\varepsilon)$ a.e. in J.

Now we must show that this n-tuple of strategies does form a BNE. Individual i cannot change the pricing function, the signal he gets, or the set K_i. By the definition of a CREE his allocation, $f_i(\varepsilon)$, gives him for a.e.

ε_i and signal observed the most preferred point in the set K^i. This holds for all the individuals and hence this n-tuple of strategies forms a BNE.

Now we shall show that any BNE of the game coincides, up to a set of measure zero, with some CREE allocation rule.

A BNE generates a pricing function $\bar{P}$ which is constructed in step 1. At step 2 we have an n-tuple of measurable strategies,

$$r_1(\varepsilon_1), ..., r_n(\varepsilon_n), \qquad \text{where} \quad r_i: J_i \to J_i.$$

These functions show what observation individual i would report as a function of his true observation.

The mapping $r(\varepsilon) = (r_1(\varepsilon_1), ..., r_n(\varepsilon_n))$ must satisfy $r(\varepsilon)$ is in J for a.e. ε in J; otherwise there is a positive probability that $|r - \varepsilon^*|$ is not zero, in which case individuals would like to announce infinitely large m_i's in order to make t_i (in the outcome function) as large as possible. Hence no choice of m_i could constitute part of an equilibrium strategy.

Let $\hat{J}$ be defined by $\hat{J} = \{\varepsilon \in J \mid r(\varepsilon) \in J\}$; since $\hat{J}$ has measure one, we may restrict our attention, for the rest of the analysis, just to elements of $\hat{J}$.

The "effective" pricing function on $\hat{J}$ is $\hat{P}(\varepsilon) = \bar{P}(r(\varepsilon))$. Following step 2 individual i's expected utility is denoted by $v_i(x_i)$, where $v_i(x_i) = E(u_i(x_i, \varepsilon) \mid \varepsilon_i, \hat{P}(\varepsilon))$.

We now break the proof into a series of claims.[11]

Claim 1. *$m^* \cdot m_i = 1$ and $\sum_{i=1}^n (y_i + w_i) = \sum_{i=1}^n w_i$, for a.e. ε in J and all $i = 1, ..., n$.*

Proof. Suppose by way of contradiction there exist an individual i and a set D of positive measure in $\hat{J}$ for which $m_i^* \cdot m_i < 1$. This would violate the BNE property of m_i. Individual i could get a preferred outcome (on a set of positive measure) by announcing a larger m_i for all ε_i in the projection of D on the ith axis. (Precise calculations justifying this are presented in Postlewaite and Wettstein [12].) Now suppose by way of contradiction there exists in $\hat{J}$ a set of positive measure for which

$$\sum_{i=1}^{n} (y_i + w_i) \neq \sum_{i=1}^{n} w_i.$$

By C's definition $\sum_{i=1}^n (y_i + w_i) \leqslant \sum_{i=1}^n w_i$ for a.e. ε; hence there must be a coordinate s and a set D in $\hat{J}$, of positive measure, for which

$$\sum_{i=1}^{n} (y_i^s + w_i^s) < \sum_{i=1}^{n} w_i^s.$$

[11] Parts of the proofs will replicate arguments that appeared in Postlewaite and Wettstein [12], and are provided in detail in order to make this presentation self contained.

Since y_i is in K^i for all $i=1, ..., n$ and for a.e. ε, we have $\hat{P} \cdot \sum_{i=1}^n (y_i + w_i) = \hat{P} \cdot \sum_{i=1}^n w_i$ for a.e. ε in J. Thus there must exist a coordinate s' and a set in $\hat{J}$ of positive measure for which $\sum_{i=1}^n (y_i^{s'} + w_i^{s'}) > \sum_{i=1}^n w_i^{s'}$, in contradiction to C's definition.

CLAIM 2. *For any ε in $\hat{J}$, the ith individual can get arbitrarily close to any net trade satisfying*:

(i) $\hat{P}(\varepsilon) \cdot z = 0$

(ii) $z + w_i \leqslant W$

(iii) $z + w_i \geqslant 0$

Proof. On observing $\hat{P}(\varepsilon)$, individual i will announce a large enough m_i and the net trade z (which is in K^i). The first part of Step 3, in the mechanism's construction, would yield $y_i = z$. The following parts would yield an arbitrarily small m^* and nullify the effects of other terms in the sum $m^* \cdot \sum_{i=1}^n m_i(y_i + w_i)$. So the final net trade will be arbitrarily close to y_i, which equals z.

Now denote the allocation rule, yielded by the mechanism at the BNE, by $f(\varepsilon)$.

CLAIM 3. *$f_i(\varepsilon)$ solves the problem*:

$$\begin{aligned} &\operatorname*{Max}_{x_i} v_i(x_i) \\ &\text{S.T.} \\ &\hat{P}(\varepsilon) \cdot x_i \leqslant \hat{P}(\varepsilon) \cdot w_i \\ &x_i \leqslant W \\ &x_i \geqslant 0 \end{aligned} \tag{1}$$

for a.e. ε in $\hat{J}$.

Proof. By the definition of K^i and the fact that $m_i^* \cdot m_i = 1$ for a.e. ε in $\hat{J}$, $f_i(\varepsilon)$ is a feasible and affordable consumption for a.e. ε in $\hat{J}$.

Assume by way of contradiction that there exists a set D of positive measure on $\hat{J}$, such that an allocation rule which is feasible for problem (1) and measurable with respect to ε_i and $\hat{P}(\varepsilon)$ yields a bundle which individual i strictly prefers to $f_i(\varepsilon)$ for all ε in D. By strict monotonicity of preferences we may assume, without loss of generality, that all those bundles are on the budget line. By claim 2 the ith individual can, by a suitable choice of strategies, get arbitrarily close to those bundles. By assumption (A3), the continuity of u_i (assumption (A2)) implies through the use of the Lebesgue Dominated Convergence Theorem that v_i is continuous as well. However, the continuity of v_i implies that the individual could achieve a preferred outcome on a set of positive measure, in contradiction to the BNE property of f.

Thus for all $i = 1, ..., n$ and a.e. in J, $f_i(\varepsilon)$ solves problem (1). Since the requirements imposed by the CREE definition must hold up to a set of measure zero, we have shown that the allocation rule yielded by the mechanism coincides a.e. in $\hat{J}$ (and hence in J) with a CREE allocation rule.

4. Concluding Remarks

We have constructed an almost continuous mechanism, implementing the SCC induced by the CREE. This mechanism performs well over a large class of preferences and greatly improves upon previously suggested mechanisms in not requiring the designer to know individual preferences.

The outcome function of the mechanism will be continuous if all the projections can be carried out continuously. This will indeed be the case if J is a convex set.

An undesirable feature of our mechanism is the non-existence of a best response in certain regions of the strategy space. This could be avoided by using discontinuous punishment schemes, but that would destroy the continuity property which we see as very important.

The signalling stage used in our mechanism was deterministic. The use of random signalling structures might be appropriate in modelling environments where individuals observe noisy information.

The most natural extension of our analysis would seem to be the introduction of production opportunities, moving away from the pure exchange setting. This however, raises several difficult issues similar in nature to those encountered in the analysis of production theory in the presence of incomplete markets.

References

1. R. J. Aumann, Agreeing to disagree, *Ann. Statist.* **4** (1976), 1236–1239.
2. L. Blume and D. Easley, Implementation of expectations equilibria, mimeo, University of Michigan, 1985.
3. J. Harsanyi, Games with incomplete information played by Bayesian players, *Manage. Sci.* **14** (1967, 1968), 159–182, 320–334, 486–502.
4. W. Hildenbrand, "Core and Equilibria of a Large Economy," Princeton Univ. Press, Princeton, NJ, 1974.
5. L. Hurwicz, Outcome functions yielding Walrasian and Lindhahl allocations at Nash equilibrium points, *Rev. Econ. Stud.* **46** (1979), 217–225.

6. L. Hurwicz, E. Maskin, and A. Postlewaite, Feasible implementation of social choice correspondences by Nash equilibria, mimeo, 1982.
7. J. Jordan and R. Radner, Rational expectations in microeconomic models: An overview, *J. Econ. Theory* **26** (1982), 201–23.
8. T. Palfrey and S. Srivastava, On Bayesian implementable allocations, *Rev. Econ. Stud.* **54** (1987), 193–208.
9. T. Palfrey and S. Srivastava, Implementation with incomplete information in exchange economies, *Econometrica* **57** (1989), 115–134.
10. A. Postlewaite and D. Schmeidler, Revelation and implementation under differential information, mimeo, University of Pennsylvania, 1983.
11. A. Postlewaite and D. Schmeidler, Implementation in differential information economies, *J. Econ. Theory* **39** (1986), 14–33.
12. A. Postlewaite and D. Wettstein, Feasible and continuous implementation of constrained Walrasian equilibria, mimeo, University of Pennsylvania, 1982.
13. D. Schmeidler, Walrasian analysis via strategic outcome functions, *Econometrica* **48** (1980), 1585–93.
14. D. Wettstein, "Implementation Theory in Economies with Incomplete Information," Ph.D. Thesis, Univ. of Minnesota, Minneapolis, 1987.