Bargaining, Coalitions and Competition*

Nir Dagan
Universitat Pompeu Fabra

Roberto Serrano
Brown University

Oscar Volij
Brown University

June 1996

*Roberto Serrano acknowledges the hospitality of Universitat Pompeu Fabra (Barcelona) and the financial support of a 1995-1996 sabbatical grant from Spain's Ministry of Education and Science. ONCE also provided research support.

Abstract

We study a decentralized matching model in a large exchange economy, in which trade takes place through non-cooperative bargaining in coalitions of finite size. Under essentially the same conditions of core equivalence, we show that the strategic equilibrium outcomes of our model coincide with the Walrasian allocations of the economy. Our method of proof exploits equivalence results between the core and Walrasian equilibria. Our model relaxes differentiability and convexity of preferences thereby covering the case of indivisible goods. (JEL classification D51, D41 and C78.
Keywords: Finite Coalitions, Strategic Bargaining, Core, Walrasian equilibrium.)

1 Introduction

The Walrasian or competitive equilibrium is the central solution concept in economics. However, from its definition it is not clear what trading procedures lead to Walrasian outcomes. In contrast to what many economists may think, the original account of the theory (Walras, 1874, pp. 83–84) does not rely on the existence of the so called Walrasian auctioneer: 'The markets which are best organized from a competitive standpoint are those in which purchases and sales are made by auctions ... Besides these markets, there are others, such as the fruit, vegetables and poultry markets, where competition, though not so well organized functions fairly effectively and satisfactorily. City streets with their stores and shops of all kinds —baker's, butcher's, grocer's, taylor's, shoemaker's, etc.— are markets where competition, though poorly organized, nevertheless operates quite adequately... '. Nonetheless, the usual formal presentation of the model includes the auctioneer implicitly due to a lack of explanation for the formation of equilibrium prices.

Negishi (1989) distinguishes two major schools in the analysis of markets. One of them considers prices as part of the economic mechanism in and out of equilibrium. This school is attributed to Cournot (1838) and Walras (1874). A second school associated with Jevons (1879) and Edgeworth (1881), attempts to consider decentralized trading mechanisms and answer the question of whether equilibrium prices will emerge as the consequence of agents' trading actions.

We can distinguish at least two major approaches in the Jevons-Edgeworth school of decentralized trading. One of them, which today is referred to as the core equivalence literature, finds its origins in Edgeworth (1881). This approach finds conditions under which core and Walrasian allocations are equivalent. If one conceives the core as a decentralized mechanism, these results give insights about the Walrasian allocations that were not provided by Walras. The shortcoming of this approach is that, although the core captures a natural idea of coalitional stability, it does not specify the trading procedure either.

A second approach models the trading procedure more explicitly and studies its strategic equilibria. As part of a recent literature, that starts with Rubinstein (1982) and Rubinstein and Wolinsky (1985) [henceforth, RW], many researchers have turned to models where a decentralized trading procedure is made explicit in a bargaining extensive form.[1] RW (1985) analyze an assignment market and claim that in a frictionless economy the strategic equilibria need not be Walrasian. This claim was challenged later by the classic paper of Gale (1986a), who constructed an alternative bargaining procedure in a continuum economy in which strategic and Walrasian equilibria coincide. Gale's work was generalized in some respects by McLennan and Sonnenschein (1991) [McLS in the sequel].

The major difference between our model and the preceeding ones is that we allow for trade to take place in coalitions with any finite number of participants, as opposed to only pairwise meetings. It turns out, though, that all Walrasian allocations can be supported by strategic equilibria in which only pairs of agents are engaged in trade. Hence, pairwise trade arises as an endogenous feature of the procedure instead of being imposed as an exogenous restriction. On the other hand, our trading procedure is indeed decentralized, since finite coalitions are negligible in the continuum. By considering multilateral meetings in which trade takes place, we are able to use the full power of the theory of the core, thereby establishing a three–fold equivalence between core, Walrasian equilibrium and non–cooperative strategic equilibrium outcomes. Our procedure resembles the Edgeworthian notion of recontracting, thus connecting naturally with the work on the non–cooperative implementation of the core in finite games and economies (see, for example, Perry and Reny (1994), Serrano (1995) and Serrano and Vohra (forthcoming)).

Our paper yields the equivalence between strategic and Walrasian equilibria under essentially the same assumptions as those needed for the core equivalence theorem. In particular, we allow for non–convex non–differentiable preferences and indivisible goods. Moreover, our assumptions also guarantee the existence of a Walrasian equilibrium, as

[1] See Osborne and Rubinstein (1990, chapter 6) and the references therein for earlier models of decentralized trade in pairwise meetings where each pair uses the Nash bargaining solution to split the gains from trade.

opposed to Gale (1986 a, b, c) and McLS (1991), which deal with open consumption sets.

Gale (1986a) lists down the conditions under which one would expect the outcome of trade to be Walrasian. These include in particular the requirement that there be a large number of individually insignificant agents. Although the core equivalence results use similar conditions, it is somewhat surprising that no use of core equivalence results appears in these decentralized strategic models. The only partial exception to this claim is the work of McLS (1991), who apply their axiomatic characterization of Walrasian allocations to the strategic analysis. This axiomatic characterization was shown by Dagan (1996) to be closely related to the replica theorem of Debreu and Scarf (1963). Our method of proof uses the core equivalence machinery. In particular, we apply the generalization of Aumann (1964) equivalence theorem due to Hammond, Kaneko and Wooders (1989) concerning the equivalence among the core, the core with respect to finite coalitions and the Walrasian allocations.

The differentiability assumption by itself that Gale (1986a, c) and McLS (1991) make is not very restrictive from an applied point of view: many models in economics incorporate it in order to allow for a closed solution of the model and for the performance of comparative statics exercises. However, the proofs of the above mentioned authors rely on additional very strong assumptions, that exclude most applied models. Gale (1986a) assumes that for each utility function the support of the endowments compatible with it is the entire consumption set. This assumption excludes the possibility of a finite type economy. Gale (1986c), who assumes a finite number of types, uses a bounded curvature assumption to prove his result, thereby excluding, for example, Cobb–Douglas utility functions on the non–negative orthant or its interior. McLS (1991) make either a bounded curvature assumption similar to Gale's (1986c) or a restriction on the equilibrium which seems to require an assumption similar to Gale's (1986a) on the primitives of the economy. In contrast, our model, which applies to very general economies, also applies to these standard cases.

We assume, like Gale (1986a, b, c), that the flow of agents enterring to the market constitutes an economy, i. e., they sum up to a finite measure. In addition, we also assume that short sales are not allowed. These two assumptions together ensure that the flow of agents out of the market is consistent with the feasibility constraint of the economy. Suppose, like McLS (1991), that the total measure of agents is finite, but short sales are allowed. In this case, nothing assures that feasibility is met: consider an arbitrary assignment of bundles to agents, and the following strategies (that do not constitute an equilibrium in McLS's game with short sales): each proposer asks for the bundle assigned arbitrarily to him and each responder accepts any proposal; agents leave the market as soon as they reach their assigned bundle. Clearly, these strategies guarantee that each agent will get with probability 1 the assigned bundle. The problem stays even if we restrict attention to the equilibria of their game. Indeed, the strategic equilibrium that McLS propose (pp. 1395–1396) to support a Walrasian equilibrium is a strategic equilibrium for any prices. That is, for an arbitrary price vector, their strategic equilibrium gives the outcome that every agent maximizes over the corresponding budget set, but the market clearing conditions may be violated. For example, the quantity of bananas people consume is larger than the supply of bananas in the economy. In light of this, one should cast doubt on the validity of such a model as a foundation of Walrasian outcomes, since the I.O.U.'s are eventually consumed by the agents. As we perceive market clearing conditions as an essential part of the Walrasian concept, we do not allow for short sales and adopt Gale's approach, which ensures feasibility in and out of equilibrium.

In the case where the sum of the measures does not constitute an economy, it is not clear what are the feasibility constraints. McLS's Theorem 2 deals with the case when the inflow and the outflow of agents have long run averages. Following this approach, we can define the flow of agents out of the market to be 'feasible' if it is consistent with the long run average of the inflow of agents. However, it is not clear to us how one can construct a model in which this kind of constraint is met in and out of equilibrium. Therefore, interpreting the result of RW (1985) as consistent with Walrasian allocations (as done by McLS's Theorem 2) is not sound: the outcome of their strategic equilibrium is consistent with a notion of feasibility, but behavior different from this equilibrium may violate the same feasibility notion.

We also address other questions left open after the work of Gale (1986a, b, c) and McLS (1991). For instance, we find in the first example of section 5 that Gale's equivalence result relies on assumptions like differentiability. In a second example in the same section, we show that the model of McLS (1991) relies crucially on the assumption of full anonymity of the trading procedure, while both our model and Gale's (1986a, b, c) do not. In particular, the same results go through whether or not agents are able to identify each other's current bundle, preferences and other characteristics that do not change over the course of the play. All the models, as noted by McLS(1991), and this applies to ours as well, require that agents not observe each others' histories.

McLS (1991) note the restrictiveness of Gale's (1986a, b, c) assumption of strictly concave utility functions in a continuum setting. One should expect that the convexifying effects of large numbers could be helpful to relax this assumption. However, we believe that McLS's solution to the problem is inadequate and provide an alternative treatment of the issue. One difference between the underlying economy and the strategic model is that in the latter the outcome (at least for an individual agent) may be random and thus preferences on random outcomes must be specified. McLS do not make any assumption regarding the concavity of utility functions; instead, they allow for short sales, which enables them to prove that outcomes of the strategic equilibria are not random. One major flaw of their approach, as already discussed, is that in a model with short sales nothing assures feasibility. In addition, a separate shortcoming of their treatment of non–convexities is that they maintain the differentiability assumption, which precludes non–convexities arising from the existence of indivisible commodities.

Gale (1986a) uses the strict concavity assumption only to ensure that the introduction of lotteries does not enlarge the set of possible utilities of the agents. Thus, what is needed is a property of risk aversion in the aggregate. We impose a condition on the quasiconcave covers of the utility functions that ensures the sufficient degree of aggregate risk aversion.

This assumption is compatible with having indivisible commodities as well as other kinds of non–convexities. Thus, our assumptions allow for a unified treatment of assignment markets à la RW (1985) and classical exchange economies à la Gale (1986a, b, c). From a methodological point of view, our assumption of aggregate risk aversion is implicit in the traditional Walrasian analysis. If it were not to hold, one should expect the emergence of markets for lotteries. We should stress that the assumption of aggregate risk aversion is sufficient to obtain Gale's results as well (within of course his restrictive subdomain of differentiable economies).

On the game theoretic side, we are careful to formalize the equilibrium notion needed to analyze a dynamic game of imperfect information. In doing so, we follow Osborne and Rubinstein (1990, chapter 8) to a large extent. These issues, missing in Gale's and McLS's work, are crucial: especially the choice of off–equilibrium beliefs, as shown in the examples of section 5.

The paper is organized as follows: Section 2 describes the underlying economic model. The non–cooperative bargaining game is described in section 3. The main result is presented in section 4 and discussed in section 5. In section 6 we show that every Walrasian allocation can be supported by a strategic equilibrium of the game, and section 7 concludes.

2 Description of the Economy

Let $(A, \mathcal{A}, \mu)$ be a measure space, where A is the set of agents, $\mathcal{A}$ is the set of measurable subsets of A and μ is an atomless measure. We denote by C the set of agents' characteristics. An element $c \in C$ is a pair $c = (u, e)$, where $u : X \to \mathbb{R}$ is a utility function, X is the consumption set and $e \in X$ is an endowment. The consumption set X is assumed to be identical for all agents and is of the form $\mathbb{R}_+^D \times \mathbb{N}^I$, where $\mathbb{N}$ is the set of non–negative integers. The consumption set includes $|D|$ divisible goods and $|I|$ indivisible goods; we assume that $|D| \geq 1$.

An economy $\mathcal{E}$ is a measurable map $\mathcal{E} : A \to C$. An allocation f is a measurable map

$f : A \to X$ that satisfies $\int_{a \in A} f(a) d\mu \leq \int_{a \in A} e(a) d\mu$. From now on, and whenever there is no danger of confusion, the domain of integration $a \in A$ will be omitted.

A coalition S can improve upon an allocation f if S has a positive measure and there exists a measurable map $g : S \to X$ such that $\int_{a \in S} g(a) d\mu \leq \int_{a \in S} e(a) d\mu$ and almost everywhere in S $u_a(g(a)) > u_a(f(a))$. The core of an economy is the set of all allocations that no coalition can improve upon.

We shall make the following assumptions on the utility functions of the agents. The assumptions can be grouped into two classes. On the one hand, we need the assumptions of Hammond, Kaneko and Wooders (1989) that guarantee the validity of their core equivalence theorem.

A0 All the commodities are present in the economy: $\int e(a) d\mu \gg 0$.

For all $c = (u, e) \in C$:

A1 The utility function u is continuous, strictly increasing in the divisible commodities, and non–decreasing in the indivisible commodities;

A2 for all $(x_D, x_I) \in I\!R^D_+ \times I\!N^I$, there exists $y_D \in I\!R^D_+$ such that $u(y_D, 0_I) > u(x_D, x_I)$; and

A3 for all $x_I \in I\!N^I, u(e) > u(0_D, x_I)$.

In addition, we need assumptions that are necessary to deal with the possibly random outcome of the strategic bargaining process.

A4 The utility functions are bounded: For all $c = (u, e) \in C$: there exists a number k_u such that $u(x) \leq k_u$ for all $x \in X$.

A5 The economy satisfies strong aggregate risk aversion (defined in the following paragraphs.)

Aggregate Risk Aversion. Gale (1986a, b, c) assumes that all individuals have strictly concave utility functions. This assumption is used in order to prove that the outcome of a strategic equilibrium of the market is a degenerate lottery.

This assumption is quite restrictive, since it excludes the possibility that the consumption sets include indivisible goods. This makes the comparison of his model with the earlier ones (e.g. RW (1985)) difficult. The work of McLS (1991) dispenses with the convexity assumption, but requires differentiability. Therefore, also in their framework indivisible goods cannot be studied. On the other hand, one should expect that the convexification effects of large numbers may help relax the assumption of individual risk aversion of preferences. These arguments motivate our assumption of aggregate risk aversion.

A lottery on allocations is a probability measure on the Borel σ–algebra of allocations. We shall say that the economy satisfies *weak aggregate risk aversion* if for every lottery L on allocations there exists an allocation g such that for almost all agents $u_a(g(a)) \geq \int_{x(a)\in X} u_a(x(a))dL$, where x denotes the allocations in the support of L.

A lottery L is *degenerate* if for almost every agent a there exists a constant k_a such that $u_a(x(a)) = k_a$ for almost all x in the support of L. The economy satisfies *strong aggregate risk aversion* if it satisfies weak aggregate risk aversion and for every lottery which is not degenerate there exists an allocation g that satisfies for almost all $a \in A$ $u_a(g(a)) \geq \int_{x(a)\in X} u_a(x(a))dL$, and there exists a set of agents of positive measure for whom the inequality is strict.

That is, aggregate risk aversion means that society cannot gain from lotteries over allocations. The strong version of the property means that if the lottery is non–degenerate, society actually will lose from having the lottery.

Although this is a property on the aggregate, we find conditions on individual preferences that are much weaker than concavity of the utility functions that imply the corresponding aggregate risk aversion properties.

First we consider an individual with a utility function $u : X \to I\!R$. We define *the* *quasiconcave cover* of u as $\hat{u} : \hat{X} \to I\!R$, where $\hat{X}$ is the convex hull of X:

$$\hat{u}(x) = max\{u(y) : x \in \hat{R}(y)\},$$

where $\hat{R}(y)$ is the convex hull of the set of all bundles that are weakly preferred to y. The assumption that the consumption set X is bounded below and assumption A1 on the utility function u ensure the existence of $\hat{u}$, as shown in Starr (1969, Appendix 3).

The first condition we consider is simply that $\hat{u}$ is concave for almost all agents. (We assume this in addition to our earlyer assumptions.)

Proposition 2.1 If the quasiconcave covers of the utility functions of almost all agents are concave, then the economy satisfies weak aggregate risk aversion.

A slightly stronger assumption on individual preferences is that the quasiconcave cover of the utility function is almost strictly concave. A function $f : X \to I\!R$ is *almost strictly concave* if for all $x_1, x_2 \in X$ such that $f(x_1) \neq f(x_2)$ and for all $\lambda \in (0, 1)$ we have that $f(\lambda x_1 + (1 - \lambda)x_2) > \lambda f(x_1) + (1 - \lambda)f(x_2)$. The difference between almost strict concavity and strict concavity is that the former requires that $f(x_1) \neq f(x_2)$ and not simply that $x_1 \neq x_2$. This difference is crucial for our purposes, because if a cover is strictly concave, the original function is strictly concave as well, whereas functions that are not even quasiconcave may have an almost strictly concave cover.

Proposition 2.2 If the quasiconcave covers of the utility functions of almost all agents are almost strictly concave, then the economy satisfies strong aggregate risk aversion.

We begin by showing the first proposition.

Proof of Proposition 2.1: Let L be a lottery on allocations. We define $EU_a(L) = \int_{x(a) \in X} u_a(x(a))dL$. We also define $E_a(L) = \int_{x(a) \in X} x(a)dL$ and $f : A \to I\!R_+^{D \cup I}$ as the function that assigns to each agent a the bundle $E_a(L)$. It is easy to see that f is integrable.

Now let $\phi(a) = \{x \in I\!R_+^{D \cup I} : \hat{u}_a(x) \geq \hat{u}_a(f(a))\}$. Clearly, $\int f(a) \in \int \phi(a)$ since it is true for every a. Let $\psi(a) = \{x \in X : u_a(x) \geq \hat{u}_a(f(a))\}$. It follows from the definition of $\hat{u}$ that $\hat{\psi}(a) = \phi(a)$, where $\hat{\psi}(a)$ is the convex hull of $\psi(a)$. Now consider $\int \psi(a)$. It follows from Liapunov's theorem that this integral is convex, and thus $\int \psi(a) = \int \phi(a)$. Therefore, $\int f(a) \in \int \psi(a)$. Thus, there exists an allocation g such that $u_a(g(a)) \geq \hat{u}_a(f(a))$ for almost all $a \in A$. It follows from the fact that all the covers $\hat{u}$ are concave that almost all agents weakly prefer the allocation g to the lottery L. $\|$

We next provide the proof of Proposition 2.2.

Proof of Proposition 2.2: Let L be a non–degenerate lottery over allocations. First it follows from Proposition 2.1 that there exists an allocation g that satisfies that $u_a(g(a)) \geq \hat{u}_a(f(a))$ for almost all $a \in A$. Since the lottery is non–degenerate, there is a positive measure of agents such that each of them is not indifferent among almost all bundles in the support of L. Now for each agent a in this set, we can have two cases:

1. There does not exist k such that for almost all x in the support of L, $\hat{u}_a(x) = k$. In this case, it follows from almost strict concavity of $\hat{u}$ that $\hat{u}_a(f(a)) > EU_a(L)$.

2. There exists k such that for almost all x in the support of L, we have that $\hat{u}_a(x) = k$. Since the agent is not indifferent among almost all bundles, there exists a positive measure of bundles in the support of L such that $\hat{u}_a(x) > u_a(x)$. Thus, $\hat{u}_a(f(a)) = k > EU_a(L)$. $\|$

Example. An example will be useful to clarify our assumptions. Consider an individual who may consume two goods: one of them is perfectly divisible and the other is indivisible, (like in an assignment market of RW (1985)). The consumer wants to consume at most one unit of the indivisible good and his reservation price in terms of the divisible commodity for the first unit of the indivisible good is 1. More formally, the commodity space is $I\!R_+ \times I\!N$. His utility function is $u(x_1, x_2) = v(x_1 + I(x_2))$ where $I(x_2) = 1$ if $x_2 \geq 1$ and $I(x_2) = 0$ otherwise, and v is a strictly increasing, strictly concave and bounded function from $I\!R_+$ to $I\!R$. Since the consumption set contains indivisible goods, there is no clear notion of risk aversion of preferences in this setting. However, the quasiconcave cover of u is almost strictly concave: $\hat{u}(x_1, x_2) = v(x_1 + min\{x_2, 1\})$. This implies for example that if there is a continuum of agents with the same preferences, the economy cannot gain from introducing lotteries over bundles.

3 Description of the Game

Time runs discretely from 1 to infinity. In each round the agents are matched at random into coalitions of finite size. At every round t there is a proportion $\alpha \in (0,1)$ that is left unmatched. For each round t and each size $n \geq 1$, there is a positive probability $p(n)$ of being matched in an n- person coalition. Thus, $p(1) = \alpha$. Matches are made randomly and for a fixed $n \geq 2$ the probability of being matched to any $n-1$ agents chosen from $n-1$ sets is proportional to the product of the measures of these sets.

When a coalition S meets, there is a 'cheap talk' phase in which every agent announces a bundle.[2] In addition, an order is chosen at random with equal probability. The first agent in the order becomes the proposer, who makes a public offer consisting of a trade $(z_j)_{j\in S}$ in which $\sum_{j\in S} z_j = 0$ and for all $j \in S, x_j + z_j \in X$ ($(x_j)_{j\in S}$ denotes the bundles held by coalition S before its meeting). Responses are also public and occur sequentially following the order. They can be one of two possible actions: 'yes,' and 'no.' The trade proposed to the coalition takes place if and only if every responder agrees to it. Every agent who is matched in round t can, if he so wishes, leave the market and consume his bundle after the bargaining session ends. Agents who are not matched in round t cannot leave the market in that round. In the next round, all agents who chose to leave abandon the market place and consume their current bundle. All other agents continue as active traders ready to be matched again.

In each round each agent recognizes the economic characteristics of the agents with

[2]The arguments in the characterization theorem are entirely independent from this phase. However, its introduction simplifies greatly the existence part, which could be complicated due to the fact that our model allows for demand correspondences.

whom he is matched. These consist of the current bundle each of them holds and their utility functions. However, they do not have information about their histories: each agent remembers only his own, but not the others'. They do not know anything about meetings that do not include them.

The restrictive information available to traders requires us to endow agents with beliefs about what happens elsewhere in the market. This must be done in order to have well defined expected utility computations. On the equilibrium path, we shall assume that these beliefs are derived from the strategies using Bayes' rule. On the other hand, off the equilibrium path we shall assume that the agents believe that the rest of the market continues to play according to the equilibrium.

The payoff to a typical trader in this market is the utility of the bundle with which he leaves the market. Thus, there is no discounting. On the other hand, if an agent never leaves the market, his utility is the utility corresponding to the zero bundle. All agents are expected utility maximizers when evaluating lotteries over bundles.

4 The Equilibrium Notion and the Main Result

A *strategic equilibrium* is a particular type of perfect Bayesian equilibrium, i.e., it is a profile of strategies, one for each agent, such that, given the beliefs specified in the previous section, every agent plays a best response to the others at every information set.

Since each agent is an entity of measure 0 in the continuum and since each of them has met only a finite number of agents in all the rounds up to round t, we can define the variable of the 'state of the market.' That is, a fixed profile of strategies played by the continuum determines the state of the market in round t as a distribution over characteristics. This happens with independence of the actions of a set of measure 0 (the history of an agent at a given point). Notice that the distribution over characteristics that we refer to as the 'state of the market' need not be supported by an allocation of the economy. Such an example can be constructed following the one found in Kannai (1970). However, a distribution is all an expected utility maximizing agent needs in order to make his calculations.

Equivalently, we could take the approach based on distributions like in Hart, Hildenbrand and Kohlberg (1974) instead of the name-based approach. We should then assume (like Gale (1986a, b, c), Osborne and Rubinstein (1990) and McLS (1991) that any two agents with identical characteristics and histories play the same strategy. This would enable us to employ the machinery developed by McLS (1991, section 3.3) in order to establish that for any given strategies the state of the market in round t in the sense of distribution of agents' characteristics is deterministic. See also Osborne and Rubinstein (1990, pp. 160-161), who show this for the finite type pure strategy case.

Now we can restate the equilibrium concept exactly as in Osborne and Rubinstein (1990), i.e., it is a perfect Bayesian equilibrium in which every agent believes both on and off the equilibrium path that the state of the market is the one that would arise if the equilibrium strategies were played. Our main result follows.

Theorem 4.1 Suppose that the economy satisfies assumptions A0–A5. In every strategic equilibrium there exists a Walrasian price p such that almost all agents leave the market in finite time with a bundle that maximizes their utility on the budget set corresponding to the price p and to their initial endowment e.

Note that an agent may leave the market receiving different bundles. However, all of them belong to the same indifference surface and maximize the agent's utility over the budget set. Also, if the maximizing bundle is unique over the budget set, then all agents with the same characteristics receive this bundle.

Proof: Consider a strategic equilibrium. All of our statements are relative to this equilibrium and to histories in which at most a set of agents of measure 0 has deviated. Since each agent's history is private information and as the matching process treats all agents alike, two agents with the same characteristics and beliefs must get the same payoff regardless of their histories. If not, the agent with the lower payoff would simply imitate the behavior of the other and get the same probability distribution over outcomes. Recall that all agents have the same beliefs about the state of the market independent of their histories.

All agents at the beginning of round t before their match has been determined believe that the state of the market corresponds to the equilibrium. Thus in the equilibrium all such agents that in addition share the same characteristics have the same expected utility. We denote this utility by $V(c,t)$. For each $c=(u,e)$, we define $w(c)=u(e)$.

Step 1: $V(c,t) \geq w(c)$ for all values c and t.

To see this, notice that every agent with characteristics c in period t can adopt the following strategy: To propose the zero trade, reject any trade, and leave the market as soon as possible. Since with probability 1 he will be matched in finite time, this strategy guarantees him a payoff of $w(c)$.

Step 2: $V(c,t) \geq V(c,t+1)$ for all values of c and t.

This assertion follows from the fact that by proposing the null trade and rejecting every offer and staying in the market, any agent in the market in round t is sure to be in the market in round $t+1$ with the same bundle as in round t.

We shall say that an agent is 'about to leave the market' if, according to the equilibrium strategies, he has already reached the situation in which his strategy tells him to leave.

Step 3: For an agent of characteristic c who is about to leave the market in round t, we have that $V(c,t+1)=w(c)$.

By step 1, we have that $V(c,t+1) \geq w(c)$. If $V(c,t+1) > w(c)$ and given that this agent is about to leave the market, he would be better off by deviating and staying in the market until round $t+1$.

Step 4: At some round t there is a positive measure of agents who are about to leave the market.

We argue by contradiction. Suppose that no positive measure of agents ever leaves the market. In this case the utility of almost all agents is that of the zero bundle. On the other hand, at any point in time there is a positive measure of agents who holds a bundle different from the zero bundle. Thereby, contradicting step 1.

Step 5: There does not exist a coalition $S \in \mathcal{A}$ with $\mu(S) > 0$, that has an S-allocation g for which $u_a(g(a)) > V(c(a), 1)$ for almost all $a \in S$, where $c(a)$ is the initial characteristics of agent a.

Assume there exist such a coalition and such an allocation. Then, by Hammond, Kaneko and Wooders (1989, Claim 1) there exists a partition of this coalition into $h+1$ coalitions $S_0, S_1, \ldots, S_h$ such that $\mu(S_1) = \mu(S_2) = \cdots = \mu(S_h) > 0$ and a list of trades $z_1, \ldots, z_h$ such that for all $a \in S_m, m = 1, \ldots, h$ $u_a(e(a)+z_m) > u_a(g(a))$ and $\sum_m z_m \leq 0$. Informally, this means that there are 'many' h-person (finite) coalitions that can improve upon their expected utility.

By step 4, at some round t a positive measure of agents is about to leave the market. We will show now that, under the contradiction hypothesis we are making, i.e., the existence of the improving coalition S, step 3 would be violated. Recall that in every round a proportion α of the agents is unmatched, and that the matching process is random. It follows from step 2 that in every round t the probability of an agent to be matched in an $h+1$-person coalition such that his h partners constitute an improving coalition is positive. Now each person who is about to leave the market can adopt the following strategy: to stay in the market and whenever being a proposer in such an improving coalition, to offer them an improving trade z which gives the proposer higher amounts of some goods without giving away any amount of the others; in all other situations, he holds on to his bundle and leaves the market at some finite date. The proposal z will be unanimously accepted by the members of the improving coalition since for each of them with characteristic c we have that:

$$V(c + z_c, t+1) \geq w(c + z_c) > V(c, t),$$

where the first inequality follows from step 1 and the second from the existence of the improving trade z for the group of responders. Clearly, this deviation gives the deviating agent a higher expected utility than the utility of his current bundle, which contradicts step 3.

Step 6: In a strategic equilibrium, there exists a Walrasian price such that almost every agent leaves the market in finite time with probability 1 holding a bundle that maximizes his utility given the budget set induced by the prices and his initial endowment.

By the previous step applied to the set of all agents A, and since the economy satisfies strong aggregate risk aversion, it follows that almost every agent receives a degenerate lottery over bundles. In addition, the resulting allocation satisfies all the core inequalities and, by Hammond, Kaneko, and Wooders (1989, Theorem 2), this allocation is Walrasian. ||

Remark: If agents are allowed to observe the state of the market every period, Theorem 4.1 and its proof go through without change when the solution concept used is the unrestricted set of perfect Bayesian equilibria.

5 Discussion

Pairwise Meetings and Non–Differentiabilities. Theorem 4.1 has shown, under very general assumptions, that all strategic equilibrium outcomes are Walrasian in a model that uses meetings of any finite size. It should be clear, though, from step 5 of the proof of Theorem 4.1, that if all gains from trade can be exhausted in coalitions of a given finite size k (e.g. assignment markets, in which $k = 2$), one can restrict the matching process to meetings of size $k + 1$.

Now we will show that restricting the matching technology to pairs may result in a friction, even in Gale's model, when preferences are not differentiable. We modify the game in the spirit of Gale, that is, if an agent trades in round t, he cannot leave the economy in that round. Suppose the off–equilibrium beliefs are allowed to be arbitrary in meetings where every agent shows up with a bundle that does not correspond to the one he should have according to the equilibrium, while beliefs are the ones used in Theorem 4.1 otherwise. It is easy to see that an analogue of Theorem 4.1 is still true for this game with these beliefs when meetings are allowed to be multilateral.

However, consider the following example. We shall first construct a perfect Bayesian equilibrium, without any restriction on off–equilibrium path beliefs, whose outcome is not Walrasian. Next we shall construct another PBE based on the previous one and with the same outcome, but where beliefs are restricted in the way described above.

There are three types of agents $i = a, b, c$ and four goods in the economy: three different types of cars (of types a, b, c) and a perfectly divisible good— money. Type i's endowment consists of one car of type i and ten dollars. Preferences are quasilinear and described as follows: each type has a valuation of 1 for its own car, a valuation of 1.5 for the car labelled with the following letter in the alphabet, and 0 for the previous letter's (of course we follow the convention that c is followed by a). If an agent holds more than one car, he derives utility only from the best car from his view-point. There are equal proportions of each type of agent in the economy. Notice that the allocation prescribed by the endowment is not Walrasian as it is not even Pareto efficient.

Consider the following strategies and beliefs that will support the endowment as a PBE: The beliefs are derived from the strategies using Bayes' rule on the equilibrium path. Off the equilibrium path, agent of type i believes that almost everyone in the market is willing to buy the car i for 1.5 dollars and sell the other cars for 1 dollar. The strategies are as follows: every responder rejects the equilibrium offer if he believes he is on the equilibrium path. Otherwise, he accepts an offer if and only if, as a result of accepting it, his income does not decrease (taking as prices the ones given by his beliefs). Every proposer offers a trade that maximizes his utility this period (in particular, the zero trade is offered on the equilibrium path). On the equilibrium path, agents leave the market after their first meeting ends; off the equilibrium path, agents leave as soon as they can after reaching their optimal consumption bundle corresponding to the budget determined by their beliefs. The outcome of this PBE is that of no trade. No proposer deviates from the zero trade because doing so would send the wrong signal to his trading partner, who will start to believe that the market is trading goods at his top ranked Walrasian equilibrium.

Next we construct a strategic equilibrium with the same outcome for the same economy. Beliefs are the ones described in the previous paragraph for the off–equilibrium behavior whenever a trader finds himself in meetings where at least one of the traders shows up with a bundle that does not correspond to the equilibrium. Otherwise, beliefs are derived from the strategies using Bayes' rule like in equilibrium. On the equilibrium path every proposer offers the zero trade and every responder accepts an offer if and only if the after-trade bundle gives him more utility than the endowment. On the other hand, in meetings where at least one of the participants shows up with a bundle different than the equilibrium one, the strategies are the ones described in the previous equilibrium.

The previous example points out that, in the presence of serious non–differentiabilities, the assumption of pairwise meetings is a friction that prevents the economy from reaching a Walrasian allocation. With differentiability, pairwise meetings do not constitute such a friction, due to the fact that every trader leaves the economy with the same marginal rate of substitution. Two interesting open questions are whether analogous examples can be constructed either when the lack of differentiability is less severe or when no conditions on preferences curvature are imposed. The difficulties there seem to be the lack of optimal plans when agents find themselves off the equilibrium path.

The Role of Anonymity in the McLS Model. Next we provide an example that shows that Proposition 2 of McLS (1991, p. 1419) is sensitive to the assumption that agents observe their trading partner's current bundle, which contradicts the authors' claim in the introduction (p. 1399, last paragraph). Thus, obtaining variants of McLS's Theorems 2 and 3 that relax the anonymity assumption is still an open question.

Consider an economy that consists of a continuum of identical agents with strictly monotone preferences over a single commodity. Recall that in McLS (1991) all meetings are in pairs and short sales are allowed. Let each agent's endowment consist of one unit of the good. The unique Walrasian allocation is that each agent consumes his endowment. We assume that agents can observe their trading partner's current bundle. Consider the following strategies that implement an outcome where each agent ends up consuming two units of the good: each proposer who has less than two units asks to receive the rest of units he is short of two. Each proposer who has two or more units proposes the zero trade. Each responder that has no more than two units accepts any trade that does not increase the proposer's bundle beyond two units, and also any trade that does not decrease his own income. Responders with more than two units accept only trades that do not decrease their incomes. Agents leave the market whenever they have two or more units (provided they can leave the market).

Notice that this example does not specify exactly the entry process, but this specification is unimportant as far as Proposition 2 of McLS goes. Proposition 2 applies to all equilibria, even to those that do not satisfy the feasibility constraints. We have not found a non–Walrasian example in which agents enterring the economy sum up to a finite measure and the market clearing conditions are met. Therefore, it is still an open question whether anonymity is required in equilibria for which the market clearing condition is satisfied.

In our model as in Gale (1986a, c), it is unimportant whether agents observe each others' bundles before they trade. For the game where agents do not observe the trading partners' current bundles, we would modify the rules of the game so that an infeasible proposal is void and the meeting ends without trade.

6 Existence

In Theorem 4.1 we have shown that all strategic equilibrium outcomes are Walrasian. Next we show the converse. That is, we find strategies that support every Walrasian outcome as a strategic equilibrium.

Theorem 6.1 Let f be a Walrasian allocation corresponding to an equilibrium price p and suppose that all agents have a maximizer over every budget set corresponding to the price vector p. Then, there exists a strategy profile that supports f as a strategic equilibrium outcome of the game.

The assumption in Theorem 6.1 is needed in our model as we can have Walrasian equilibria with some prices equal to 0. In this case, it could be that some agents (that constitute a set of measure 0) do not have an optimal bundle in their budget sets even in the equilibrium allocation. On the other hand, the definition of a strategic equilibrium is very demanding, in two respects: (1) restrictions to behavior are placed in and out of equilibrium, and (2) every agent must play a best response to the others. While Walrasian equilibrium is usually defined in this context when almost every agent is maximizing over the budget set, such a definition is not possible for a perfect equilibrium in an extensive form game. Notice for example that the deviation argument in step 5 of Theorem 4.1 need no longer hold since the improving coalition (a set of measure 0) need not be playing a best response to the deviating proposal. Of course, if equilibrium prices are all positive, or the model is a standard assignment market, (e.g. like in RW (1985)), the assumption in Theorem 6.1 is satisfied. In the models of Gale (1986a, b, c) and McLS (1991), such an assumption is also needed since they work with open consumption sets.

Proof: Consider the following strategy profile: For each agent a let $h(a, e)$ be a function that assigns to each agent a with holdings e a bundle from the agents' demand correspondence with respect to the price p and the income pe. The function h is restricted so that $h(a, e) = f(a)$ if $pe = pe(a)$. In all meetings every trader announces during the 'cheap talk' phase the bundle assigned to him by the selection defined above. In multilateral meetings (those with at least three agents), the proposer offers the 0 trade; in bilateral meetings, the proposer offers a trade according to the trading rule defined below (which is based on Gale (1986b)). As in Gale (1986b), we shall distinguish between the behavior of one of the divisible commodities (say, commodity 1) and that of the other $|D| + |I| - 1$ goods. While according to the trading rule, an agent's excess demand in commodities other than 1 is non–increasing, commodity 1 serves to balance the budget whenever there is no pure coincidence of wants.

We shall denote the proposer by a_0 and the responder by a_1. Subscripts denote agents and superscripts denote commodities. Let z_0 be $h(a_0, e(a_0, t)) - e(a_0, t)$, where $e(a_0, t)$ are the holdings of a_0 in round t, and let z_1 be the trade announced by the responder in the 'cheap talk' phase. Define the set $B(z_0, z_1)$ as the set of vectors $x \in I\!R^D \times \mathbf{Z}^I$ satisfying the following conditions (where $\mathbf{Z}$ denotes the set of integers).:

(i)$|x^h| \leq |z_i^h|, i = 0, 1; h \geq 2$

(ii) $0 \leq (-1)^i x^h z_i^h, i = 0, 1; h \geq 2$

(iii) $e(a_i) + (-1)^i x \in I\!R_+^D \times I\!N^I$ and $px = 0$.

If the net trade x is proposed and accepted, the proposer's new endowment $e(a_0, t+1) = e(a_0, t) + x$ and the responder's $e(a_1, t+1) = e(a_1, t) - x$. The trade proposed is denoted by

$$g(z_0, z_1) = \text{arg.max}\{-\sum_{h \geq 2} \exp(-|x^h|) : x \in B(z_0, z_1)\}$$

.

In all meetings, every responder a who currently holds e that did not achieve the bundle $h(a, e)$ accepts a trade if and only if his income (the value of his holdings evaluated at the prices p) does not decrease. If he already achieved the bundle $h(a, e)$, then he accepts if and only if his income increases. Every agent leaves the market as soon as he achieves the bundle $h(a, e)$.

Denote by $u^*(c, p)$ the maximum utility that an agent with characteristics c can achieve over the budget set determined by his endowment and the prices p. We will show that if every agent behaves according to the specified strategies, almost every agent of characteristic c achieves a bundle (corresponding to the allocation f) that yields $u^*(c, p)$ in finite time.

Notice first, as in Gale (1986b), that if agents follow the specified strategies, it is not possible for an agent to increase his income as evaluated by the prices p. Next we will show that an agent a ends up at the bundle $f(a)$ in finite time with probability 1. This will show that the proposed strategies are a strategic equilibrium. That is, given that there is no way to increase one's income, the proposed strategies induce a random path that takes each agent to his chosen maximizer $f(a)$ over the budget set determined by e and p.

For each agent a and for prices p define the excess demand as follows (for convenience and given that the Walrasian prices p are fixed throughout, we shall drop p from the expressions below): $\phi(a,t) = f(a) - e(a,t)$, where $e(a,t)$ are the holdings of agent a at round t. Notice that, given the strategies specified above, every agent travels along the frontier of his budget set which means that $f(a)$ continues to be a utility maximizer for agent a.

As we said above, we shall distinguish between the behavior of commodity 1 and that of the others. The trading rule is constructed so that the absolute value of the excess demand of every agent in all goods but 1 does not increase. On the other hand, good 1 serves to balance the budget whenever there is no pure coincidence of wants. Thus we define for each agent a the following number: $\beta(a,t) = \sum_{h\geq 2} |\phi^h(a,t)|$. That is, for each agent a the statistic $\beta(a,t)$ indicates the sum of absolute values of excess demands in all goods but 1 in round t. We will next show that as time goes on the distribution of $\beta(a,t)$ converges weakly in measure to a degenerate distribution on 0.

Recall that μ denotes the measure of characteristics in the economy, i.e., the measure of characteristics of all agents who are active in the market plus that of the agents who already left the market. The random matching process and the specified strategies lead to new distributions of characteristics at every round, and hence to new distributions of the statistic $\beta(a,t)$. We shall concentrate on an arbitrary path determined by a particular realization of the different random variables at play (the coalitions that meet and the roles of each agent in each meeting). We show then that, along this path, the distribution of $\beta(a,t)$ converges in measure to the degenerate distribution on 0.

The space of characteristics C at each round t is the Cartesian product of a fixed space of utility functions with $\mathbb{R}^D_+ \times \mathbb{N}^I$. The evolution of the economy is thus described by a sequence of measurable maps from the set of agents A to the set C. Given that the set of utility functions is fixed throughout the model, any Cauchy sequence of such measurable maps must converge to a measurable map from A to C. To see this, notice that, after having fixed the utility functions, the marginal of characteristics c on agents' endowments e allows us to consider a Cauchy sequence of integrable maps from A to $\mathbb{R}^D_+ \times \mathbb{N}^I$. Endowing this space of integrable maps with the topology induced by the supremum norm, it is easy to see that such a sequence converges to an integrable map into $\mathbb{R}^D_+ \times \mathbb{N}^I$ as this is a complete space. Notice that the marginal of the measure μ_t on utility functions is constant. We can then abuse notation slightly and denote by μ_t round t's measure on the agents' endowments and not on characteristics. Thus, using Hildenbrand (1974, p.50) the sequence of measures $\{\mu_t\}$ is tight and has a convergent subsequence to μ^*. Without loss of generality, suppose the sequence itself converges to μ^*.

By the properties of the trading rule, g we must have that for a given constant $\tau > 0$, $\int_{\beta(a,t)\geq\tau} \mu(a,t)$ converges to 0 as time goes to infinity. To see this, we argue by contradiction. Suppose that the limiting measure μ^* is not the one concentrated at 0. Since μ^* is the limiting measure, it must be the case that the measure of agents trading positive amounts of goods when the distribution of $\beta(a,t)$ is approximately μ^* must be arbitrarily close to 0: For all $\epsilon > 0$ there exists a T such that for all $t > T$ we have that $\int[\mu_t - \mu^*] < \epsilon$. If μ^* is not the one concentrated at 0, there must exist a positive measure of agents whose characteristics satisfy that $\phi^h(c) > 0$ for some good h. By Walras' law which holds at each step of this time path, there must also exist a positive fraction of agents whose characteristics satisfy that $\phi^h(c) < 0$. Since the matching process is random, there exists a positive probability that agents in these two situations will meet. Finally, given the trading rule, these agents will trade at least in good h, which is a contradiction, i.e., there exists $\epsilon > 0$ such that for all T there exists $t > T$ with $\int[\mu_t - \mu^*] > \epsilon$ as $g(z_0, z_1)$ stays bounded away from 0 for a positive fraction of meetings.

As for convergence in finite time, the arguments are identical to those in Gale (1986b, section 7). $\|$

Remark: If we were taking the approach based on distributions, where we are forced to assign the same strategy to all the agents of the same type we need to modify the proposed strategies in the following way: In the first round of trade, each agent randomizes over the different bundles assigned to this type in the underlying Walrasian allocation. Then, this bundle is treated as $f(a)$ in our proof.

7 Conclusion

By using the insights of the core equivalence theorem, this paper has presented a model of decentralized trade through bargaining in coalitions of finite size. This has allowed us to obtain equivalence results among core, Walrasian and strategic equilibrium allocations for a wide class of large exchange economies, including non–convex non–differentiable preferences and indivisible goods.

References

Aumann, R. J. (1964), Markets with a Continuum of Traders, *Econometrica* 32, 39–50.

Cournot, A. (1838), *Researches into the Mathematical Principles of the Theory of Wealth.* English edition (ed. N. Bacon) Macmillan, 1897.

Dagan, N. (1996), Recontracting and Competition, Economics working paper 152, Universitat Pompeu Fabra, Barcelona.

Debreu, G. and H. Scarf (1963), A Limit Theorem on the Core of an Economy, *International Economic Review* 4, 235–246.

Edgeworth, F. Y. (1881), *Mathematical Psychics*, Kegan Paul Publishers, London.

Gale, D., (1986a), Bargaining and Competition Part I: Characterization, *Econometrica,* 54, 785–806.

Gale, D., (1986b), Bargaining and Competition Part II: Existence, *Econometrica,* 54, 807–818.

Gale, D., (1986c), A Simple Characterization of Bargaining Equilibrium in a Large Market without the Assumption of Dispersed Characteristics, working paper 86–05, CARESS, University of Pennsylvania, Philladelphia, PA.

Hammond, P., M. Kaneko and M. H. Wooders (1989), Continuum Economies with Finite Coalitions: Core, Equilibria and Widespread Externalities, *Journal of Economic Theory* 49, 113–134.

Hart, S., W. Hildenbrand and E. Kohlberg (1974), On Equilibrium Allocations as Distributions on the Commodity Space, *Journal of Mathematical Economics* 1, 159–166.

Hildenbrand, W. (1974), *Core and Equilibria of a Large Economy*, Princeton University Press, Princeton, NJ.

Jevons, W. S. (1879), *The Theory of Political Economy*, Mcmillan, London.

Kannai, Y. (1970), Continuity Properties of the Core of a Market, *Econometrica* 38, 791–815.

McLennan, A. and H. Sonnenschein (1991), Sequential Bargaining as a Foundation of Competitive Equilibrium, *Econometrica* 59, 1395–1424.

Negishi, T. (1989), *History of Economic Theory*, North Holland, Amsterdam.

Osborne, M. J. and A. Rubinstein, (1990), *Bargaining and Markets.* Academic Press, San Diego, CA.

Perry, M. and P. J. Reny (1994), A Non–Cooperative View of Coalition Formation and the Core, *Econometrica* 62, 795–817.

Rubinstein, A., (1982), Perfect Equilibrium in a Bargaining Model, *Econometrica*, 50, 97–109.

Rubinstein, A. and A. Wolinsky, (1985), Equilibrium in a Market with Sequential Bargaining, *Econometrica*, 53, 1133–1150.

Serrano, R. (1995), A Market to Implement the Core, *Journal of Economic Theory* 67, 285–294.

Serrano, R. and R. Vohra (forthcoming), Non–Cooperative Implementation of the Core, *Social Choice and Welfare.*

Starr, R. (1969), Quasi–Equilibria in Markets with Non– Convex Preferences, *Econometrica* 37, 25–38.

Walras, L. (1874), *Elements of Pure Economics, or the Theory of Social Wealth*, English edition (ed. W. Jaffé), Reprinted by Orion Editions, Philadelphia, PA, 1984.