Markets Versus Negotiations: the Predominance of Centralized Markets*

Zvika Neeman† and Nir Vulkan‡

December, 2003

Abstract

The paper considers the consequences of competition between two widely used exchange mechanisms, a "decentralized bargaining" market, and a "centralized" market. In every period, members of a large heterogenous group of privately-informed traders who each wish to buy or sell one unit of some homogenous good may opt for trading through one exchange mechanism. Traders may also postpone their trade to a future period. It is shown that trade outside the centralized market completely unravels. In every perfect equilibrium, almost all trade takes place in the centralized market. Almost no trade (measure zero) ever occurs through direct negotiations.

JEL CLASSIFICATION NUMBERS: C78, D40, L10.

KEYWORDS: centralized markets, decentralized markets, decentralized bargaining, market microstructure, competition.

---

* Acknowledgement to be added.

† Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, Center for Rationality and Economics Department, Hebrew University of Jerusalem, Israel 91904; Email `zvika@BU.edu`, http://people.bu.edu/zvika.

‡ The Said Business School, Park End Street, Oxford OX1 1HP, UK. Email `Nir.Vulkan@sbs.ox.ac.uk`, http://www.worc.ox.ac.uk/fellows/vulkan/.

1. Introduction

This paper considers the consequences of competition between two widely used exchange mechanisms, a "decentralized bargaining" market, and a "centralized" market. Competition assumes the following form: in every period, surviving and new members of a large heterogenous group of privately-informed traders who each wish to buy or sell one unit of some homogenous good may opt for trading through either (1) direct negotiations with other buyers and sellers (a decentralized bargaining market), or (2) a centralized market mechanism. If they so wish, traders may also postpone their trade to a future period.

A *decentralized bargaining market* is an idealization of what takes place in a bazaar, or a Middle-Eastern Suq. Buyers and sellers are matched with each other and bargain over the terms of trade. If an agreement is reached, the traders leave the market. Otherwise, they return to the general pool of traders, are possibly re-matched with another trader, and so on. Importantly, transaction prices in such a market typically vary across the different matches depending, among other things, on the individual traders' costs and willingness to pay. Consequently, decentralized bargaining is characterized by the fact that different traders may transact at *different* prices at the same time.[1] In contrast, a *centralized market* is a form of exchange with a single price and centralized clearing. It is characterized by the fact that in any point in time, all those traders who transact do so at the *same* price. Examples include the undergraduate textbook definition of a competitive market, a sealed-bid double-auction, and a *call market*.[2]

The study of the outcome of such a competition is of interest for two reasons. First, the question of what form of exchange is likely to attract large volumes of trade is an important theoretical and practical problem. To solve this complex problem, it is not enough to analyze the properties of different exchange mechanisms in isolation. Because traders' choices of where to trade are endogenous, the very existence of a competing exchange mechanism may affect the outcome in any given mechanism. In other words, the question is what kind of

[1] Under "extreme" circumstances (e.g., in a stationary environment with anonymous, infinitely patient, traders, and with no aggregate uncertainty), decentralized bargaining may give rise to trade in only one (Walrasian) price (Gale, 1987). Such circumstances are excluded by our assumptions. See the discussion below in Section 3.2.

[2] Call markets are used, among other things, to determine the daily opening prices of the stocks listed in the New York Stock Exchange, and to fix copper and gold price in London (Schwartz, 1988).

exchange mechanisms will flourish when traders are free to choose the exchange mechanism through which to transact. Obviously, because exchange mechanisms that may initially be attractive to sellers may be shunned by buyers and vice-versa, mechanisms that generate large volumes of trade must be sufficiently attractive to *both* buyers and sellers. Second, although the description of the two competing exchange mechanisms is extremely stylized, the comparison is between a "traditional" and a "modern" form of exchange. Understanding the forces that determine the consequences of such a competition may shed some light on the development of actual market mechanisms.

The main result of this paper is that, under fairly general assumptions, in every perfect equilibrium, almost all trade is conducted through the centralized market – almost no opportunities for mutually advantageous trade exist outside the centralized marketplace (here and in the rest of the paper, "almost all" should be understood as "except possibly for a set of Lebesgue measure zero"). Obviously, because traders cannot trade alone, there also exists a (non-perfect) equilibrium in which all traders trade only through direct negotiations, however, our result implies that such an equilibrium is "unstable." If in some period a small measure of traders were accidentally forced into the centralized market, then the rest would want to follow, undermining the equilibrium.

The approach in this paper is distinguished by the fact that, in contrast to standard models that impose assumptions about traders' *behavior* and then derive the implications of these assumptions with respect to market structure, here the assumptions are imposed directly on the distribution of transaction prices under the two competing mechanisms. This "reduced-form" approach allows us to bypass the main difficulty that is associated with the standard approach, namely the characterization of equilibrium properties, which seems to be intractable in any but the simplest models. In contrast, the approach followed in this paper permits the consideration of a general dynamic setup with discounting, heterogeneity, asymmetric information, and aggregate uncertainty.

There is a vast literature on the microstructure of markets and trading institutions. This literature can be divided into several broad categories. First, there is a large literature that has confined its attention to the analysis of different market mechanisms in isolation. In this literature, comparisons between different market mechanisms are usually done from the perspective of the seller, asking which mechanism a single seller would prefer under the assumption that buyers have no choice but to participate in the chosen mechanism (as in, e.g., Milgrom and Weber, 1982). A second category, into which this paper belongs, consists of papers that consider the case in which traders choose through which one of a small number of given mechanisms to conduct their trades (see, e.g., Gehrig, 1993; Rust and Hall, 2003, and the references therein).[3] Finally, there exists a voluminous related literature in finance, which emphasizes the importance of transaction costs, information, adverse selection, and transparency, but which pays less attention to the strategic issues considered here (for a recent survey of this literature, see Madhavan, 2000).

The paper that is most closely related to ours is Rust and Hall (2003) who consider the case in which buyers and sellers can choose between trading through "middlemen" or a "market maker." As we show below, trading through middlemen, or as they are sometimes called, dealers or brokers, is a particular example of what we call direct negotiations. Trading through the market maker in Rust and Hall's paper is similar to trading through the centralized market in our model, except that the monopolistic market maker in Rust and Hall (2003) charges a positive bid-ask spread, which introduces a small wedge between the price paid by buyers and the price received by sellers. The pattern of trade described by Rust and Hall (2003) is similar to the one described in this paper. Buyers with high willingness to pay and sellers with low costs trade through the market maker, while others resort to searching and trading through the decentralized "search" market. However, while in our model the unraveling of the decentralized search market is complete, in Rust and Hall's model, the wedge that the market maker introduces between the price paid by buyers and received by sellers prevents their search market from completely unraveling.

The rest of the paper proceeds as follows. In the next section, we describe a simple example that illustrates our basic argument and insight. In Section 3, we present the general model and the details of modelling direct negotiations and centralized markets. Analysis of the model is presented in Section 4, and a few concluding remarks are offered in Section 5.

2. An Example

We describe a simple static example that provides an intuition for our main result. The general model is dynamic, has a large number of heterogenous traders, and unlike the example

[3] A number of papers permit traders to choose from a large number of possible trade mechanisms (see, e.g., McAfee, 1993; Peters, 1994; and subsequent literature) but in models where competing sellers choose a type of auction through which to sell and buyers select in which seller's auction to participate.

that assumes specific forms of direct negotiations and centralized markets, is consistent with many different ways of modelling direct negotiations and centralized markets.

Consider an environment with six traders, three buyers and three sellers. Each buyer wants to buy, and each seller wants to sell, one unit of some homogenous good. The buyers' willingness to pay for the good are 10, 9, and 2, respectively; and the sellers' costs of producing the good are 0, 1, and 8, respectively. Note that the static nature of the example implies that there is no reason to abstain from trade.

If all the traders trade in a centralized market where they behave as price-takers, then any price $p \in [2, 8]$ can serve as a market clearing price. Suppose, for simplicity, that the price that prevails in the market is $p = 5$. The two buyer's types with the willingness to pay of 10 and 9 trade with the two seller's types whose costs are 0 and 1. The payoff to the buyer whose willingness to pay is 10 and to the seller whose cost is 0, is 5; and the payoff to the buyer whose willingness to pay is 9 and to the seller whose cost is 1, is 4. The other buyer's and seller's types do not trade in the market, and obtain, each, a payoff of 0.

Suppose on the other hand that the traders engage in direct negotiations with each other. Suppose further that this negotiation assumes the following form: a first stage of random matching between the buyers and sellers, followed by a second stage of split-the-surplus bargaining. In this case, the expected payoff to the buyer whose willingness to pay is 10 is 3.5 because with probability $\frac{1}{3}$, the buyer is matched with the seller whose cost is 8, trades at the price 9, and obtains a payoff of 1, with probability $\frac{1}{3}$, the buyer is matched with the seller whose cost is 1, trades at the price $\frac{11}{2}$, and obtains a payoff of $\frac{9}{2}$, and with probability $\frac{1}{3}$, the buyer is matched with the seller whose cost is 0, trades at the price 5, and obtains a payoff of 5. Similarly, the expected payoff to the buyer whose willingness to pay is 9 is 3, and the expected payoff to the buyer whose willingness to pay is 2 is $\frac{1}{2}$ because when this buyer is matched with the seller whose cost is 8, no trade takes place. Similarly, the expected payoff to the seller whose cost is 8 is $\frac{1}{2}$ and the expected payoffs to the sellers whose costs are 0 and 1 are 3.5 and 3, respectively.

Obviously, the two buyer's types with the high willingness to pay and the two seller's types with the low costs (the "weak" types) are better off in the centralized market compared to direct negotiations. Even if they alone switch to trading through the centralized market, they are still better off as they can still trade at the competitive equilibrium price $p = 5$. However, once they switch, the remaining buyer and seller become worse off since they lose their ability to trade. They, too, may switch to the centralized market, but this will not improve their situation, since they do not get to trade in the centralized market either.

Intuitively, what makes the centralized market more attractive to the two buyer's types with the high willingness to pay and the two seller's types with the low cost is that, relative to direct negotiations, the extent to which their high willingness to pay and low costs are translated into higher and lower prices, respectively, is smaller. Consequently, the weak types of the buyer and seller are led into trading in the centralized market, which in turn, leads to the unraveling of trade through direct negotiations.

In this simple example, it is easy to imagine a direct negotiation procedure that would lead to exactly the same outcome as the centralized market (simply change the matching process to one that ensures that the buyer with willingness to pay 10 is matched with the seller whose cost is 0, and the buyer with willingness to pay 9 is matched with the seller whose cost is 1). Obviously, in this case we cannot obtain our result that almost all trade must take place in the centralized market. The literature on decentralized bargaining (surveyed in Osborne and Rubinstein, 1990, and discussed in more detail in the next section) has devoted much attention to the question of how likely is "frictionless" decentralized bargaining to give rise to the centralized market (Walrasian) outcome. In any case, the assumptions imposed below preclude this possibility.

Finally, it is important to emphasize that what makes the centralized market more attractive to the weak traders' types is not just the fact that it offers a potentially more efficient form of matching than direct negotiations, but also the way in which the surplus from trade is distributed among the different traders' types. More specifically, the distribution of surplus in the centralized market is biased in favor of weak traders' types, which is what starts the process of unraveling. This can be best seen by comparing the expected payoff to different traders' types under the centralized market and under direct negotiations with efficient matching. Suppose for example that direct negotiation assumes the following form: a first stage of matching in which with probability $\frac{1}{2}$ the buyers with willingness to pay 10, 9, and 2 are matched with the sellers whose costs are 0, 1, and 8, respectively; and with probability $\frac{1}{2}$ the buyers with willingness to pay 10, 9, and 2 are matched with the sellers whose costs are 1, 0, and 8, respectively. Under this direct negotiation procedure, the expected payoffs to the buyers with willingness to pay 10, 9, and 0, and to the sellers with costs 0, 1, and 8, are 4.75, 4.25, and 0, respectively. The buyer with willingness to pay 9 and the seller with cost 1 are better off with this type of direct negotiations than under the centralized market, but the buyer with willingness to pay 10 and the seller with cost 0, or the weakest types of traders, are worse off, and once these types switch into trading in the centralized market, other traders' types would be better off following them there.

3. The Model

We consider a dynamic model with a large number (a continuum) of buyers and sellers of some discrete homogenous good. Time is also discrete and is indexed by $t \in \{1, 2, ...\}$. In each period, each seller has at most one unit to sell, and each buyer is interested in buying at most one unit.[4] It is commonly known that sellers' costs and buyers' willingness to pay range from 0 to 1. Traders are characterized by their types: their willingness to pay for one unit of the good if buyers, and their cost of producing one unit if sellers. "Weak" buyer types have high willingness to pay, and "weak" seller types have low costs.

We assume that a positive mass of new traders appears in every period. The realized cumulative distributions of the new buyers' and sellers' types in every period are assumed to be increasing, differentiable, and to be equal to zero at zero. Thus, in every period t, a positive measure of buyers and sellers is present. In addition to the new traders, the group of traders at t may also include traders that have appeared in the previous period but who for some reason did not trade. The realized cumulative distributions of buyers' and sellers' types in every period are therefore increasing, differentiable, and equal to zero at zero. Every trader knows his own type but may be uncertain about the measure of other traders and the distribution of their types. No restrictions are imposed on the distribution of the size of the incoming group of traders in any period (except that it has a positive and finite support) and on the joint distribution of traders' types. In particular, traders' types need not be independently distributed, and the realized distribution of traders' types in any given period need not be deterministic. That is, the model admits aggregate uncertainty.

Every period, present buyers and sellers may either attempt to trade through a centralized market or through direct negotiations. Traders may also refrain from trade and wait for the next period. A trader who for some reason does not trade in any given period re-appears in

[4] This implies no loss of generality compared with the assumption that traders are each interested in trading a finite number of units of the good. Traders that are interested in trading $k < \infty$ units of the good can be treated as k different traders.

the next period (with the same type) with probability $0 \leq \delta \leq 1$. The parameter δ may also be interpreted as the traders' discount factor. Buyers and sellers are assumed to be (risk neutral) expected utility maximizers. A buyer with a willingness to pay v who transacts at the price p obtains the payoff $v - p$, and a seller with cost c who transacts at the price p obtains the payoff $\delta^{\tau}(p - c)$. Traders who disappear without trading, or who never trade, obtain the payoff zero.

In every period $t \geq 1$, the traders' choices about whether to attempt to trade through the centralized market or through direct negotiations may depend on their own type, and on the history of trade in the centralized market and direct negotiations, respectively, in the periods prior to t. Obviously, traders' choices may also depend on their beliefs about what other traders, "new" and "old," will do. A Nash equilibrium is a sequence of profiles of traders' choices of if and where to trade, where each trader's choice is optimal given other traders' choices. To simplify the analysis, we assume that traders who are indifferent between trading in the centralized market and direct negotiation in some period t, opt for trading, if at all, through direct negotiations.

We describe the details of trade in centralized markets and direct negotiations below.

3.1. Centralized Markets

For our purposes, a centralized market may be idealized as follows: in every period $t \in \{1, 2, ...\}$ in which the market operates, a market price p_t^M is determined and is taken as given.[5] Ex-ante, the traders, who may not know the realized number (measure) of other traders and the distribution of other traders' types, may perceive the market price as stochastic. We assume that almost all the traders share the same beliefs about the distribution of the market price, which we denote by Q_t.[6] In particular, p_t^M may be given by the Walrasian or market-clearing price (which is obtained at the intersection point of the traders' true non-increasing demand and non-decreasing supply curves).

We assume that in every period in which the market operates, every buyer with willingness to pay $v \geq p_t^M$ and every seller with cost $c \leq p_t^M$ who opted for trading in the centralized

[5] Depending on the way in which the market price is determined, the existence of a market price may require that at least one trader opts for trading in the centralized market.

[6] This can be consistent with the presence of aggregate uncertainty if the latter is due to uncertainty about the realized measure of traders in the market.

market trade at the price p_t^M and obtain a payoff of $v - p_t^M$, and $p_t^M - c$, respectively. Other traders who opted for trading in the market do not trade and reappear in the next period with probability δ.[7] Thus, for almost all $c, v \in [0,1]$, the expected payoff to a seller whose cost is c from participating in the centralized market in period t is given by

$$\begin{aligned} S_t^M(c) &= E\left[\max\left\{p_t^M - c, 0\right\}\right] \\ &= \int_c^1 (p-c)\, dQ_t(p) \end{aligned}$$

and the expected payoff to a buyer whose willingness to pay is v from participating in the centralized market in period t is given by

$$\begin{aligned} B_t^M(v) &= E\left[\left\{v - p_t^M, 0\right\}\right] \\ &= \int_0^v (v-p)\, dQ_t(p)\,. \end{aligned}$$

Note that both of these payoff functions are continuous and monotone decreasing or increasing, respectively. Another property of $S_t^M(c)$ and $B_t^M(v)$ of which we make use in the proof below is the following. In every period in which the centralized market operates, for almost every $c \in [0,1]$,

$$\begin{aligned} S_t^M(c) &= \int_c^1 (p-c)\, dQ_t(p) \\ &\geq \int_0^1 (p-c)\, dQ_t(p) \\ &= E\left[p_t^M\right] - c. \end{aligned}$$

Similarly, for almost every $v \in [0,1]$,

$$\begin{aligned} B_t^M(v) &= \int_0^v (v-p)\, dQ_t(p) \\ &\geq v - E\left[p_t^M\right]. \end{aligned}$$

Together, these two inequalities imply that

$$v - c \leq B_t^M(v) + S_t^M(c) \tag{1}$$

[7] Thus, if p_t happens not to be a market clearing price at t, then in case of excess demand, a sufficiently large inventory from which all demands can be satisfied is implicitly assumed to exist, and in case of excess supply, it is implicitly assumed that a sufficiently large inventory may be created to allow all sellers to sell. In this case, an additional restriction that requires the size of the expected inventory to be bounded from above and below may be added but is not required for our results.

for almost every $v, c \in [0, 1]$.

In case the price p_t^M is equal to the market-clearing price at t, this "reduced-form" description of a centralized market is consistent with the more specific descriptions of competitive markets with privately informed traders. For example, Rustichini, Satterthwaite, and Williams (1994) idealize a *double-auction* market as follows: buyers and sellers simultaneously submit bids and offers which are aggregated to form demand and supply functions. The crossing of the demand and supply functions determines an interval $[a, b]$ of possible market-clearing prices. Trade occurs at the price $p = \frac{a+b}{2}$ among the buyers who bid at least p and the sellers whose offers were no more than p. In case a shortage or a surplus exists, the allocation is carried out as far as possible by assigning priority to sellers whose offers were the smallest and buyers whose bids were the largest. If this does not complete the allocation, then a fair lottery determines which of the remaining traders on the long side of the market trade.[8] Since in large double-auction markets competition among the traders almost completely eliminates the ability of any single trader to affect the price, traders have strong incentives to behave almost as "price-takers" in the sense of bidding approximately their true costs and willingness to pay. (When there is a continuum of buyers and sellers, single traders cannot affect the market price, and so have a *dominant* strategy to bid their true willingness to pay and cost, respectively.) Under the assumption that the buyers' and sellers' willingness to pay and costs, respectively, are independently generated from some given distributions (as well as a few other assumptions), Rustichini et al. confirm that price-taking behavior is also plausible when the number of traders is not very large, and moreover, traders have private information with respect to their costs and willingness to pay. They report the results of simulations that show that the convergence to "truthful" bidding already occurs when there are as few as three traders on each side of the market. Since the focus of our analysis is on the market price, another relevant consequence of Rustichini et al. analysis is that, under their assumptions, the market price converges to the market clearing price that would prevail in the market if the traders revealed their willingness to pay and costs truthfully (pp. 1049-1050).

More generally, Gul and Postlewaite (1992) establish general conditions, interpreted as conditions about the degree of "information smallness," under which competitive equilibria (with market-clearing prices) are robust to the introduction of asymmetric information about

[8] As noted by Rustichini et al. this trading mechanism resembles a call market.

preferences. The relevant conditions are satisfied here.

Finally, we have assumed that traders may choose to stay out of the centralized market at any given period, but that if they opt into the centralized market, then they behave competitively. To the extent that traders may behave non competitively in the centralized market, for example by refusing to trade if the price is does not turn out to be sufficiently favorable, then their expected payoffs in the centralized market would be even higher, which would further strengthen our result about the impossibility of trade outside the centralized market.

3.2. Direct Negotiation

As in the case of centralized markets, we also adopt a reduced form approach to model the process of direct negotiations among the traders. For every willingness to pay v and cost c, let $\mu_t^B(v)$ and $\mu_t^S(c)$ denote the measure of buyers and sellers with willingness to pay and costs smaller or equal to v and c, respectively, who opt for trading through direct negotiations in period t. We assume that in every period traders who opted to participate in direct negotiations are matched into pairs of one buyer and one seller according to a density function $f_t(v, c)$.[9] Conditional on being matched, a buyer and seller with types v and c, respectively, trade with each other with probability $x_t(v, c)$, at the expected price $p_t^N(v, c)$. It is natural to assume that trade is measure preserving, or that the measure of units sold is equal to the one bought, but for our purposes it is sufficient to assume that it cannot be that a positive measure of buyers trades with a positive probability with a null set of sellers, or vice-versa.

In addition, we assume that:

1. For every $t \geq 1$, $f_t(\cdot, \cdot)$, $x_t(\cdot, \cdot)$, and $p_t^N(\cdot, \cdot)$, are continuous functions. They may

[9] That is, conditional on being matched with a positive probability, a buyer with willingness to pay v is matched with a seller with cost $c \in [c', c'']$ with probability

$$\frac{\int_{c'}^{c''} f_t(v, c)\, d\mu_t^S(c)}{\int_0^1 f_t(v, c)\, d\mu_t^S(c)},$$

and a seller with cost c is matched with a buyer with willingness to pay $v \in [v', v'']$ with probability

$$\frac{\int_{v'}^{v''} f_t(v, c)\, d\mu_t^B(v)}{\int_0^1 f_t(v, c)\, d\mu_t^B(v)}.$$

depend on the history of trade in direct negotiations and the centralized market, as well as on the distributions of traders' types in direct negotiations and in the centralized market in period t itself. In particular, the expected transaction price $p_t^N(v,c)$ may be "low" if the ratio between the measures of buyers and sellers who opted for direct negotiations at t is large, and "high" if the ratio is small.

2. For every $t \geq 1$, and buyer's and seller's types $v, c \in [0,1]$,

$$c \leq p_t^N(v,c) \leq v.$$

That is, trader prefer not to trade rather than trade at a price that generates a negative payoff.

3. The price function $p_t^N(v,c)$ is nondecreasing in v and c, and strictly increasing in either v or c.

The last property captures the intuition that exactly because of their "weakness," weak buyer types (i.e., buyers with a high willingness to pay) are likely to pay relatively higher prices and weak seller types (seller with low costs) are likely to accept relatively lower prices. This property is satisfied in many models of bargaining including Nash's (1950) model of axiomatic bargaining, Rubinstein's (1982) model of alternating offers bargaining, and Myerson and Satterthwaite's (1983) optimal mechanism for bilateral bargaining under asymmetric information. This is the property that distinguishes our model of direct negotiations from our model of a centralized market, where individual traders are assumed to have no effect on the price at which they transact, and consequently, weaker types are not relatively disadvantaged because of their weakness.

Given $\mu_t^B(v)$, $\mu_t^S(c)$, $f_t(\cdot,\cdot)$, $x_t(\cdot,\cdot)$, and $p_t^N(\cdot,\cdot)$, denote the expected payoffs conditional on trade from engaging in direct negotiations in period t of the buyers and sellers by $B_t^{N|trade}(v)$ and $S_t^{N|trade}(c)$, respectively.[10] Our assumptions about $f_t(\cdot,\cdot)$, $x_t(\cdot,\cdot)$, and $p_t^N(\cdot,\cdot)$, imply that the two functions, $B_t^{N|trade}(v)$ and $S_t^{N|trade}(c)$, are continuous, and satisfy the following property, which we employ repeatedly in the proof below.

[10] That is, $B_t^{N|trade}$ is given by

$$B_t^{N|trade}(v) = \frac{\int_0^1 \left(v - p_t^N(v,c)\right) x_t(v,c) f_t(v,c)\, d\mu_t^S(c)}{\int_0^1 x_t(v,c) f_t(v,c)\, d\mu_t^S(c)}$$

if buyer v trades with a positive probability ($\int_0^1 x_t(v,c) f_t(v,c)\, d\mu_t^S(c) > 0$) and is assumed to be equal to

Lemma 1. *For any cost and willingness to pay $0 \leq c^* < v^* \leq 1$, if almost all the sellers with costs $c \geq c^*$ who opt for direct negotiations at t trade almost surely with buyers with willingness to pay $v \leq v^*$, and almost all the buyers with willingness to pay $v \leq v^*$ who opt for direct negotiations at t trade almost surely with sellers with costs $c \geq c^*$, then*

$$B_t^{N|trade}(v) + S_t^{N|trade}(c) < v - c, \tag{2}$$

for almost every cost $c \geq c^*$ and willingness to pay $v \leq v^*$ that are sufficiently close to c^* and v^*, respectively.

Proof. If almost no buyer with willingness to pay $v \leq v^*$ trades through direct negotiations with a positive probability at t, then, by assumption, almost no seller with cost $c \geq c^*$ may trade through direct negotiations with a positive probability at t either, and vice-versa. Because in this case almost all buyers and sellers trade through direct negotiations with probability zero, $B_t^{N|trade}(v) = S_t^{N|trade}(c) = 0$ for almost every willingness to pay $v \leq v^*$ and cost $c \geq c^*$, and the conclusion of the lemma follows.

If a positive measure of buyers with willingness to pay $v \leq v^*$ trade through direct negotiations with a positive probability at t, then, by assumption, also a positive measure of sellers with costs $c \geq c^*$ must trade through direct negotiations at t with a positive probability, and vice-versa. We say that a seller's type, c', is *isolated* if there exists an open set of sellers' types, O, such that $O \cap \{c \in [0,1] : c \text{ opts for direct negotiation}\} = \{c'\}$. The measure of isolated sellers' types is zero, therefore, by assumption, only a measure zero of buyers may trade through direct negotiations with a positive probability with isolated sellers' types. Hence, almost all buyers who trade through direct negotiations trade with sellers' types who are not isolated. Similarly, almost all sellers who trade through direct negotiations trade with buyers' types who are not isolated. Continuity of $f_t(\cdot,\cdot)$ and $x_t(\cdot,\cdot)$ therefore implies that for almost every v, c, $v \leq v^*$, $c \geq c^*$, who opt for direct negotiations at t, and for every $\varepsilon > 0$,

$$\Pr\left(v \text{ trades with } \widehat{c} \in \left(\widehat{\widehat{c}}, \widehat{\widehat{c}} + \varepsilon\right)\right) > 0,$$

zero otherwise, and $S_t^{N|trade}$ is given by

$$S_t^{N|trade}(c) = \frac{\int_0^1 \left(p_t^N(v,c) - c\right) x_t(v,c) f_t(v,c)\, d\mu_t^B(v)}{\int_0^1 x_t(v,c) f_t(v,c)\, d\mu_t^B(v)}$$

if seller c trades with a positive probability and is assumed to equal to zero otherwise.

for some $\widehat{\widehat{c}} \geq c^*$, and

$$\Pr\left(c \text{ trades with } \widehat{v} \in \left(\widehat{\widehat{v}} - \varepsilon, \widehat{\widehat{v}}\right)\right) > 0.$$

for some $\widehat{\widehat{v}} \leq v^*$.

The monotonicity of the price function $p_t^N(\cdot, \cdot)$ then implies that

$$B_t^{N|trade}(v) \leq v - p_t^N(v, c^*),$$

for almost every $v \leq v^*$, because conditional on trade, almost every buyer v pays a price that is almost surely larger than $p_t^N(v, c^*)$. A similar argument implies that,

$$S_t^{N|trade}(c) \leq p_t^N(v^*, c) - c,$$

for almost every $c \geq c^*$. Moreover, the fact that $p_t^N(\cdot, \cdot)$ is strictly increasing in v or c implies that at least one of the previous two inequalities is strict. Hence, adding these two inequalities, it follows that

$$B_t^{N|trade}(v) + S_t^{N|trade}(c) < v - c + p_t^N(v^*, c) - p_t^N(v, c^*)$$

for almost every $v \leq v^*$ and $c \geq c^*$. In particular,

$$B_t^{N|trade}(v^*) + S_t^{N|trade}(c^*) < v^* - c^*. \tag{3}$$

Now, because both $B_t^{N|trade}(v) + S_t^{N|trade}(c)$ and $v - c$ are continuous in v and c, and for $v = v^*$ and $c = c^*$, inequality (3) holds, it must also be the case that inequality (3) holds for almost all $c \geq c^*$ and $v \leq v^*$ that are sufficiently close to c^* and v^*, respectively. ■

Note that the conclusion of Lemma 1 is not satisfied if the price in a bilateral transaction is fixed independently of the willingness to pay and cost of the buyer and seller, as would be the case if the price were exogenously fixed, or in a large centralized market.

As indicated above, the restrictions on $f_t(\cdot, \cdot)$, and $x_t(\cdot, \cdot)$, and especially on the price function $p_t^N(\cdot, \cdot)$, are consistent with many forms of direct negotiations. We describe two examples of such direct negotiation procedures below.

Example 1 (Direct Negotiations). Fix a sequence of real numbers $\{\alpha_t\}_{t \in \{1,2...\}}$ such that for every t, $\alpha_t \in (0, 1)$. Suppose that in every period t the procedure of direct negotiations between the buyers and sellers assumes the form of random matching into pairs of one buyer and one seller, followed by split-the-surplus bilateral bargaining where buyers capture a fraction $\alpha_t \in (0,1)$ of the available surplus, and the sellers get the rest. When a buyer with willingness to pay v and a seller with cost c are matched in period t, they transact at the price $\alpha_t c + (1-\alpha_t)\,v$ if $v \geq c$ and refrain from trade otherwise.

Example 2 (Direct Negotiations). The direct negotiation procedure resembles the one describes in the previous example, except that each period is divided into $n \geq 1$ sub-periods. Traders who were matched with partners with whom they could not profitably trade, are re-matched in the next sub-period, and the procedure repeats itself, as long as the period is not over and some traders that have not yet traded remain. Traders discount the payoffs obtained at the k-th sub-period, $0 \leq k \leq n$, at the rate $1 - \frac{k}{n}(1-\delta)$.

Example 3 (A Dealers' market). This example is based on Spulber (1996). A dealer market consists of continuums of measure one of heterogenous buyers, sellers, and middlemen. Traders and middlemen discount future profits at a rate $\delta < 1$, and in every period, each trader exits the market with an exogenously specified probability $\lambda > 0$. The initial distribution of buyers' and sellers' types is the uniform distribution on $[0,1]$, and whenever a buyer or seller trades or exists the market, he is replaced by another buyer or seller whose type is drawn from the uniform distribution on $[0,1]$. The only way for buyers and sellers to trade is through middlemen who quote bid and ask prices. The middlemen are infinitely lived and each set a pair of stationary bid and ask prices to maximize their expected discounted profits. The middlemen are each characterized by their transaction costs which are uniformly distributed over the interval $[0,1]$. A middlemen with transaction cost (type) $k \in [0,1]$ sets bid and ask prices $a(k)$ and $b(k)$, respectively. Buyers and sellers engage in sequential search. Each period, a searcher obtains a single price quote from one, randomly drawn, middleman. It can be shown that this market has a unique stationary equilibrium. In this equilibrium, bid and ask prices are uniformly distributed over some interval. The fact that in this equilibrium buyers with a higher willingness to pay and sellers with lower costs are willing to buy from middlemen who quote higher and lower ask and bid prices, respectively, implies that the implied matching and price functions $f_t(\cdot,\cdot)$, $x_t(\cdot,\cdot)$, and $p_t^N(\cdot,\cdot)$, satisfy the assumptions described above.

Example 4 (An Auctions' market). In every period t, each seller sells his object through an auction, and buyers choose in which seller's auction to participate. Suppose that sellers' may each specify a reserve price, and buyers choose randomly among sellers who specified the same reserve price. The fact that in equilibrium, in all standard auctions (English, Dutch, first-price, second-price, all-pay) buyers with high willingness to pay pay more in expectation, and that sellers' optimal reservation prices are nondecreasing in their costs, implies that the assumptions above are satisfied.

A large literature has analyzed the conditions under which "frictionless" decentralized bargaining may give rise to the Walrasian outcome, in which all buyers with willingness to pay above the Walrasian price and all sellers with costs below the Walrasian price trade at the Walrasian price (for a survey of this literature, see Osborne and Rubinstein, 1990). This literature has shown that in environments with infinitely patient and anonymous traders and with no aggregate uncertainty, the Walrasian outcome may prevail (see, e.g., Gale, 1986, 1987). The Walrasian outcome is obviously incompatible with our assumptions because it is incompatible with the conclusion of Lemma 1.[11] However, if decentralized bargaining is not "frictionless," or more specifically, if traders are *not* infinitely patient (that is, traders' discount factor δ is strictly less than 1), then the conclusion of this literature is that the Walrasian outcome is impossible. Therefore, the fact that our assumptions about direct negotiations preclude the Walrasian outcome need not imply great loss of generality.[12]

4. Equilibrium Analysis

In this section, we show that in any period in which the centralized market operates, almost all the "serious" traders, namely those traders' types that would trade with a positive probability in the centralized market, who prefer to trade in that period, opt for trading in the centralized market. Consequently, almost any trade that occurs, occurs in the centralized market. Almost no trade takes place through direct negotiations. It follows that only two types of equilibria may exist: one in which all trade occurs through the centralized market, and another where the centralized market never operates and all trade occurs through direct negotiations. As we demonstrate below, while the former type of equilibrium is *perfect*, the

[11] For the same reason, a direct negotiation procedure in which in every period, trade is taking place at the expected price in the centralized market in the same period, is also inconsistent with our assumptions.

[12] Obviously, to the extent that decentralized bargaining does give rise to the Walrasian outcome, then the centralized market, which is generally believed to produce the Walrasian outcome, still "dominates" decentralized bargaining in the sense that what happens in the former determines the outcome in the latter.

latter is not. That is, while the first type of equilibrium is robust to "trembles," or uncertainty about traders' choices, the latter is not. The former type of equilibrium strategies is therefore more "reasonable [...] for rational and intelligent players to choose in this game" (Myerson, 1991, p. 213).

The idea of the proof is to show that in any period in which the centralized market operates, trade through direct negotiations "unravels" as traders with relatively weak types express their preference for trading through the centralized market.

Because a trader who fails to trade at t may trade at a future period, it does not necessarily follow that traders would prefer to trade through the exchange mechanism that provides the higher expected payoff at t.[13] However, as the next lemma shows, traders would prefer the centralized market if their (unconditional) expected payoff there is larger than their expected payoff conditional on trade in direct negotiations.

Lemma 2. *For every period $t \in \{1, 2, ...\}$ and for every buyer with willingness to pay v who opts for trading in period t: if $B_t^M(v) > B_t^{N|trade}(v)$, then the buyer strictly prefers to trade through the centralized market than to trade through direct negotiations at t. Similarly, for a seller with cost c, if $S_t^M(c) > S_t^{N|trade}(c)$, then the seller strictly prefers to trade through the centralized market than to engage in direct negotiations at t.*

Proof. We prove the lemma for buyers. The proof for sellers is similar. We introduce the

[13] Suppose for example that in every period, in the centralized market a trader trades with probability $\frac{2}{3}$ and obtains an expected payoff conditional on trade of $\frac{1}{2}$, and in direct negotiations, the trader trades with probability $\frac{1}{9}$ and obtains an expected payoff conditional on trade of $\frac{9}{10}$. Although in any given period the expected payoff from trading in the centralized market is higher $\left(\frac{1}{3} > \frac{1}{10}\right)$, a patient trader would maximize his payoff by repeatedly trying to trade though direct negotiations.

following notation: let

$B_t^{M|trade}(v)$ = the expected payoff conditional on trade to a buyer whose willingness to pay is v from participating in the market in period t;

$B_t^N(v)$ = the expected payoff to a buyer with willingness to pay v from engaging in direct negotiations in period t;

$P_t^M(v)$ = the probability that a buyer with a willingness to pay v trades in the market in period t;

$P_t^N(v)$ = the probability that a buyer with a willingness to pay v trades in direct negotiations in period t;

$B_{t\to\infty}(v)$ = the expected discounted payoff to a buyer with willingness to pay v in period t who chooses optimally whether and where to trade.

For any period t, a buyer with willingness to pay v strictly prefers to trade through the centralized market than to engage in direct negotiations if and only if his expected discounted payoff from doing so is higher, or

$$P_t^M(v)B_t^{M|trade}(v) + \left(1 - P_t^M(v)\right)\delta B_{t+1\to\infty}(v) > P_t^N(v)B_t^{N|trade}(v) + \left(1 - P_t^N(v)\right)\delta B_{t+1\to\infty}(v). \quad (4)$$

Suppose that

$$B_t^M(v) > B_t^{N|trade}(v).$$

Because $B_t^M(v) = P_t^M(v)B_t^{M|trade}(v)$, $0 \le P_t^M(v) \le 1$, and by assumption $B_t^{M|trade}(v) \ge 0$,

$$B_t^{M|trade}(v) \ge B_t^M(v).$$

We may also assume that

$$B_t^{N|trade}(v) = \frac{B_t^N(v)}{P_t^N(v)} \ge B_t^N(v).$$

Otherwise, $B_t^{N|trade}(v) < B_t^N(v) < 0$ and the buyer will certainly not engage in direct negotiations at t. It therefore follows that both

$$B_t^{M|trade}(v) > B_t^{N|trade}(v)$$

and

$$B_t^M(v) > B_t^N(v).$$

Because unless $B_t^{M|trade}(v) \geq \delta B_{t+1\to\infty}(v)$ the buyer is better off refraining from trade in period t, we may also assume that $B_t^{M|trade}(v) > \delta B_{t+1\to\infty}(v)$. (Note that if $B_t^{M|trade}(v) = \delta B_{t+1\to\infty}(v)$, then $B_t^{M|trade}(v) > B_t^{N|trade}(v)$ implies that the buyer is strictly better off trading through the market at t.) Inequality (4) is equivalent to:

$$P_t^M(v) > P_t^N(v)\frac{B_t^{N|trade}(v) - \delta B_{t+1\to\infty}(v)}{B_t^{M|trade}(v) - \delta B_{t+1\to\infty}(v)}. \tag{5}$$

Because $P_t^M(v)B_t^{M|trade}(v) = B_t^M(v)$ and $P_t^N(v)B_t^{N|trade}(v) = B_t^N(v)$, (5) may be rewritten as,

$$P_t^M(v) < P_t^N(v) + \frac{B_t^M(v) - B_t^N(v)}{\delta B_{t\to\infty}(v)}. \tag{6}$$

Finally, if $P_t^M(v) \geq P_t^N(v)$, then $B_t^{M|trade}(v) > B_t^{N|trade}(v)$ implies (5) and so (4). And if $P_t^M(v) \leq P_t^N(v)$, then $B_t^M(v) > B_t^N(v)$ implies (6) and so again (4). The conclusion of the lemma follows. ■

Every Nash equilibrium sequence of distributions of buyers' and sellers' types in the centralized market and direct negotiations, respectively, induces a sequence of continuous functions $\{B_t^M\}_{t\geq 1}$, $\{S_t^M\}_{t\geq 1}$, $\left\{B_t^{N|trade}\right\}_{t\geq 1}$, and $\left\{S_t^{N|trade}\right\}_{t\geq 1}$. Given these sequences of functions define two sequences $\{v_t^*\}_{t\geq 1}$ and $\{c_t^*\}_{t\geq 1}$ as follows: for every $t \geq 1$,

$$v_t^* \equiv \min\left\{v \in [0,1] : B_t^M(v') > B_t^{N|trade}(v') \text{ for a.s. every } v' > v\right\},$$

and

$$c_t^* \equiv \max\left\{c \in [0,1] : S_t^M(c') > S_t^{N|trade}(c') \text{ for a.s. every } c' < c\right\}.$$

The definitions of v_t^* and c_t^* imply that the set of buyers and sellers with types larger than v_t^* and smaller than c_t^*, respectively, who opt for direct negotiations at t is null.[14] Note that continuity of the functions B_t^M, S_t^M, $B_t^{N|trade}$, and $S_t^{N|trade}$ ensures that both v_t^* and c_t^* are well defined. It also ensures that $B_t^M(v) \geq B_t^{N|trade}(v)$ for every $v > v_t^*$ and $S_t^M(c) \geq S_t^{N|trade}(c)$ for every $c < c_t^*$. If $B_t^M(v) \leq B_t^{N|trade}(v)$ for almost every v close to one, then $v_t^* = 1$, and

[14]The numbers v_t^* and c_t^* are the *essential supremum* and the *essential infimum*, respectively, of the sets of buyers and sellers who opt for direct negotiations at t.

if $S_t^M(c) \leq S_t^{N|trade}(c)$ for almost every c close to zero, then $c_t^* = 0$. Finally, in periods t in which $v_t^* < 1$ or $c_t^* > 0$, $B_t^M(v_t^*) = B_t^{N|trade}(v_t^*)$ or $S_t^M(c_t^*) = S_t^{N|trade}(c_t^*)$, respectively.

The previous lemma implies that almost every buyer with willingness to pay $v > v_t^*$ and almost every seller with cost $c < c_t^*$ strictly prefer to trade through the centralized market than trough direct negotiations in period t, respectively. We show that for every $t \in \{1, 2, ...\}$ in which the centralized market operates, $0 \leq v_t^* \leq c_t^* \leq 1$. Because in every period t in which the centralized market operates almost all the buyers and sellers that engage in direct negotiations have willingness to pay smaller or equal than v_t^*, and costs larger or equal than c_t^*, respectively, the inequality $v_t^* \leq c_t^*$ implies that almost no opportunities for mutually beneficial trade exist outside the centralized marketplace at t.

Lemma 3. *In every Nash equilibrium, in every period* $t \in \{1, 2, ...\}$ *in which the centralized market operates,* $0 \leq v_t^* \leq c_t^* \leq 1$.

Proof. Fix some Nash equilibrium. Fix some $t \in \{1, 2, ...\}$ in which the centralized market operates. Suppose that $c_t^* < v_t^*$. We show that this implies a contradiction. The definitions of c_t^* and v_t^* imply that almost all the sellers with costs $c \geq c_t^*$ who opt for direct negotiations at t trade almost surely with buyers with willingness to pay $v \leq v_t^*$, and almost all the buyers with willingness to pay $v \leq v_t^*$ who opt for direct negotiations at t trade almost surely with sellers with costs $c \geq c_t^*$. It therefore follows that for almost every $v \leq v_t^*$ and $c \geq c_t^*$ that are close to v_t^* and c_t^*, respectively,

$$\begin{aligned} v - c &= \left(v - E\left[p_t^M\right]\right) + \left(E\left[p_t^M\right] - c\right) \\ &\leq B_t^M(v) + S_t^M(c) \\ &\leq B_t^{N|trade}(v) + S_t^{N|trade}(c) \\ &< v - c, \end{aligned}$$

where the first inequality follows from (1), the second inequality follows from the definitions of c_t^* and v_t^*, and the third inequality follows from Lemma 1. A contradiction. ■

We summarize our results in the following proposition.

Proposition. *In every Nash equilibrium, in every period in which the centralized market operates, almost all those buyers and sellers that trade, trade through the centralized market.*

Almost no trade occurs through direct negotiations in periods in which the centralized market operates.[15]

The proposition asserts that in periods in which the centralized market operates almost all trade takes place in the centralized market. The result is for *almost all* as opposed to *all* trade because we cannot rule out the possibility that a measure zero of buyers and sellers trade with each other through direct negotiations at an expected price that is equal to the expected centralized market price in the same period. Continuity of the functions $f_t(\cdot,\cdot)$ and $x_t(\cdot,\cdot)$ implies that not only is it the case that trade through direct negotiations is of measure zero, but that it can only take place between pairs of "isolated" buyers and sellers, that is buyers' and sellers' types that are surrounded by a neighborhood where no other buyer's or seller's type, respectively, opts for direct negotiations.

The intuition for this result is the following. Inequality (1) may be interpreted as implying that centralized markets are characterized by the fact that high value traders keep the entire additional marginal surplus generated by their type. In contrast, in direct negotiations, as Lemma 1 shows, relatively weak types of traders are forced to share the surplus they generate with others. This causes relatively weak types of traders to prefer the centralized market which causes direct negotiations to unravel.

Specifically, fix for some period t distributions of buyers' and sellers' types in the centralized market and direct negotiations, respectively. Suppose that the induced functions B_t^M, S_t^M, $B_t^{N|trade}$, and $S_t^{N|trade}$, give rise to the upper and lower bounds, $v_t^* < 1$ and $c_t^* > 0$, respectively. These functions and bounds are depicted in the following figure.

[15] The multiplicity of Nash equilibria is a consequence of the fact in any period, a measure zero of traders may opt for trading through direct negotiations. Also, traders may "coordinate" on declining to trade through both mechanisms in certain periods.

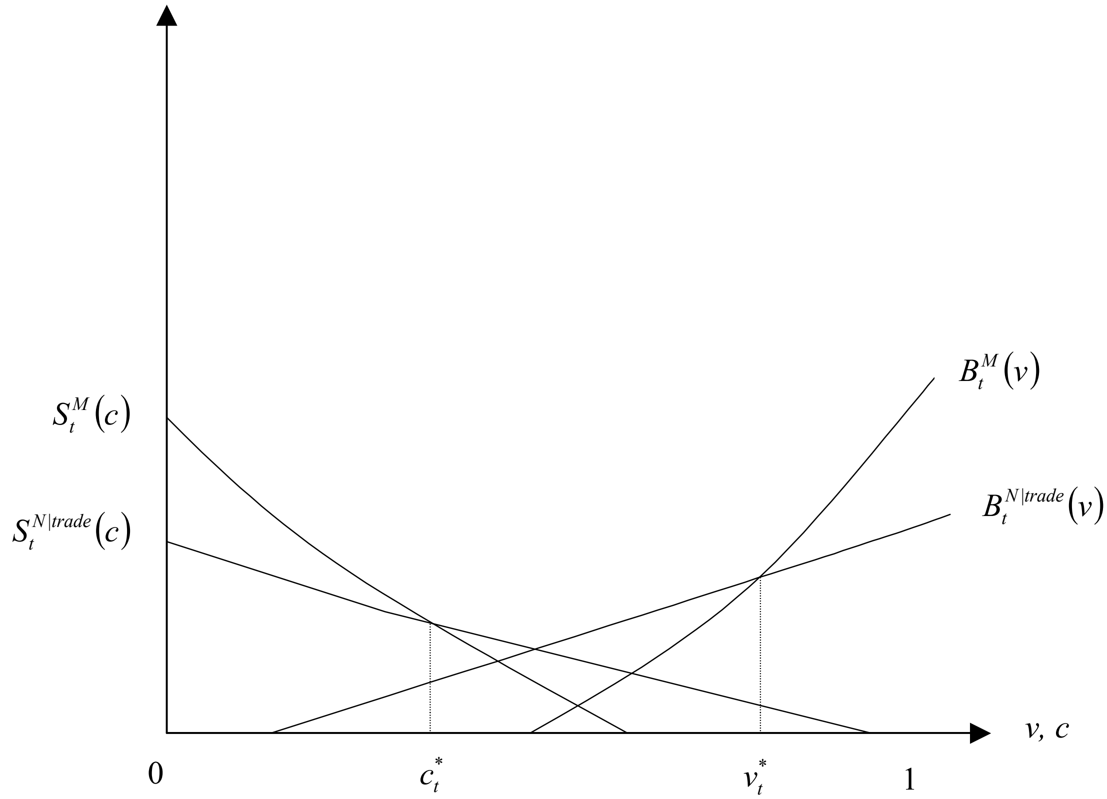

Figure 1: B_t^M, S_t^M, $B_t^{N|trade}$, $S_t^{N|trade}$, v_t^*, and c_t^*

As can be seen in the figure, buyers with willingness to pay above v_t^* and sellers with costs below c_t^*, who according to the original distributions opted for direct negotiations, would switch to the centralized market. Their switching would change the distributions of buyers' and sellers' types in the centralized market and direct negotiations, respectively. The functions B_t^M, S_t^M, $B_t^{M|trade}$, and $S_t^{M|trade}$, and the bounds v_t^* and c_t^* may now be recomputed. The expected price in the centralized market and so also B_t^M and S_t^M may not change much. But, because it was the buyers with relatively high willingness to pay and sellers with relatively low cost (the weak types) who switched to the centralized market, the distributions of buyers' and sellers' types in direct negotiations will be more concentrated on buyer types with low willingness to pay and seller types with high cost. Consequently, the expected payoffs conditional on trade for those traders who still opt for direct negotiations will be lower, the newly computed v_t^* will be lower than the previously computed one, and the newly computed c_t^* will be higher than the previously computed one. The process of unraveling will now continue as buyers with willingness to pay above the newly computed

v_t^* and sellers with cost below the newly computed c_t^* will switch to the centralized market and so on, until the computed v_t^* and c_t^* will be such that $0 \leq v_t^* \leq c_t^* \leq 1$.[16]

A *perfect* equilibrium (Selten, 1975) is a refinement of Nash equilibrium that requires that traders' strategies be robust against trembles or small uncertainty about other traders' strategies (Myerson, 1991, p. 216). The equilibrium in which the centralized market never operates and all trade is conducted through direct negotiations is obviously not perfect. If at some period t the traders would "tremble" by, say, choosing to trade through the centralized market, direct negotiations, or not to trade at all with probability at least ε for some small $\varepsilon > 0$, then a large enough number of traders would opt for trading in the centralized market, the centralized market would operate, and the lemmas above would apply. It follows that engaging in direct negotiations cannot possibly be a best response to such a profile of strategies and the equilibrium where every trader adopts such a strategy cannot be perfect. The next corollary follows:

Corollary. *In every perfect equilibrium, almost all those buyers and sellers that trade, trade through the centralized market. Almost no trade ever occurs through direct negotiations.*[17]

Proof. The previous analysis established the existence of only two types of equilibria. One in which all trade is conducted through direct negotiations and the other where almost all trade is conducted through the centralized market. The preceding discussion showed that the former type of equilibrium is not perfect, since a perfect equilibrium exists (see, e.g., Myerson, 1991, p. 221), it must be of the second type.[18] ■

[16] For simplicity, we ignored in the discussion above the possibility of trade between pairs of isolated traders. It can incorporated into the argument by re-plotting the graphs of $B_t^{N|trade}$ and $S_t^{N|trade}$ so that they meet the graphs of B_t^M and S_t^M from below in "isolated" points to the right of v_t^* and to the left of c_t^*, respectively.

[17] As in the Proposition, the multiplicity of equilibria is a consequence of the fact that a measure zero of traders may opt for trading through direct negotiations and that traders may "coordinate" on declining to trade through both mechanism in certain periods.

[18] Two comments:

(1) The existence theorem we refer to is for finite games whereas the game presented here is infinite. It has to be verified that the existence proof carries over to our case.

(2) It is possible to directly prove that the equilibrium in which almost all traders opt for trading in the centralized market whenever they trade is perfect. Trembling implies that the centralized market operates in every period and the lemmas imply that in this case, for any trader who prefers to trade in any given

Similar arguments show that the equilibria where all trade is conducted through the centralized market are also proper, persistent, and stable (Myerson, 1991, ch. 5) while the equilibrium where all trade is conducted through direct negotiations is neither.

5. Concluding Remarks

The argument presented here is that when faced with the choice, buyers and sellers will opt for trading through a centralized market over engaging in (some form of) direct negotiations. Nevertheless, some transactions, even in homogenous goods, are still conducted through direct negotiations. A number of possible explanations may be given for this. We discuss these explanations in the context of the model described in this paper.

First, it may be that the traded good is not really homogenous. Problems associated with quality and credibility may arise, and traders may prefer the relative security of establishing long term trading relationships with a small number of trustworthy trading partners, where the prospect of engaging in future trade serves as a disciplinary device against opportunistic behavior, to trading in an anonymous centralized market where relatively little protection against opportunistic behavior is provided (Kranton, 1996).[19]

Second, participation in a centralized market may entail some costs that we have not taken into account here (transportation costs and the fact that some markets convene only infrequently are possible examples). However, if these costs are similar to those incurred under direct negotiations, our main result should still hold.

Third, ensuring the constant operation of the centralized market, which is necessary for our main result, is a public good, or more precisely, a public service. A centralized market may not dominate other forms of exchange if no one is willing to assume the responsibility for the orderly provision of this public service. As shown by Rust and Hall (2003), if the centralized market is organized by a market maker who charges a positive bid-ask spread, then the unraveling of direct negotiations will not be complete.

Fourth, we have assumed in our analysis that traders are risk neutral. Another reason to prefer direct negotiations over a centralized market is that the former may allow *risk averse* traders to reduce their exposure to the centralized market's volatility by directly

period, opting for trade in the centralized market is optimal.

[19] However, even in Kranton's (1996) model, all trade will eventually be conducted through the centralized market if its initial size is sufficiently large.

negotiating to trade at the *expected* centralized market price. However, because more risk averse or pessimistic traders should also be willing to pay to reduce their exposure to risk, an argument similar to the one presented in this paper implies that a centralized *futures* market that insures against this volatility will again dominate decentralized private mutual insurance agreements.

Finally, even when, say, because of the presence of transaction costs in the centralized market, the unraveling of trade outside the centralized market does not go all the way towards eliminating trade through direct negotiations, our model still provides an insight about the relative willingness of different types of traders to trade through different forms of exchange. Centralized markets are characterized by the fact that high value traders keep the entire additional marginal surplus generated by their type. In contrast, in many models of negotiations (Nash, Rubinstein, Myerson-Satterthwaite) traders are forced to share this surplus with others. This causes weak types of traders to prefer centralized markets which would cause direct negotiations to unravel, wholly or partially.

References

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

Gale, D. (1987) "Limit Theorems for Markets with Sequential Bargaining," *Journal of Economic Theory 43, 20-54.*

Gehrig, T. (1993) "Intermediation in search markets," *Journal of Economics and Management Strategy* 2, 97-120.

Gul, F. and A. Postlewaite (1992) "Asymptotic Efficiency in Large Exchange Economies with Asymmetric Information," *Econometrica* 60, 1273-1292.

Kranton, R. (1996) "Reciprocal exchange: a self-sustaining system," *American Economic Review* 86, 830-851.

Madhavan, A. (2000) "Market microstructure: A Survey," *Journal of Financial Markets* 3, 205-258.

McAfee, P. R. (1993) "Mechanism design by competing sellers," *Econometrica* 61, 1281-1312.

Milgrom, P. and R. Weber (1982) "A theory of auctions and competitive bidding," *Econometrica* 50, 1089-1122.

Myerson, R. (1991) *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, Massachusetts.

Myerson, R. and M. A. Satterthwaite (1983) "Efficient mechanisms for bilateral trading," *Journal of Economic Theory* 29, 265-281.

Nash, J. F. (1950) "The bargaining problem," *Econometrica* 18, 155-162.

Osborne, M. and A. Rubinstein (1990) *Bargaining and Markets*, San Diego, London, Sydney, and Toronto: Harcourt Brace Jovanovich and Academic Press.

Peters, M. (1994) "Equilibrium mechanisms in a decentralized market," *Journal of Economic Theory* 64, 390-423.

Royden, H. L. (1988) *Real Analysis,* New York: Macmillan.

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.

Rust, J. and G. Hall (2003) "Middlemen versus market makers: a theory of competitive exchange," *Journal of Political Economy* 111, 353-403.

Rustichini, A., M. A. Satterthwaite, and S. R. Williams (1994) "Convergence to efficiency in a simple market with incomplete information," *Econometrica* 62, 1041-1063.

Schwartz, R. A. (1988) *Equity Markets: Structure, Trading, and Performance*, New York: Harper & Rowe.

Selten, R. (1975) "Reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory* 4, 25-55.

Spulber, D. (1996) "Market making by price-setting firms," *The Review of Economic Studies* 63, 559-580.