Automated Negotiation in Many-to-Many Markets for Imperfectly Substitutable Goods[+]

Christ Preist, Carlos Mérida-Campos[1]
Trusted E-Services Laboratory
HP Laboratories Bristol
HPL-2002-197
August 4th, 2003*

E-mail: cwp@hplb.hpl.hp.com, dmerida@lsi.upc.es

In this paper, we present an agent which is able to negotiate the buying and selling of imperfectly substitutable goods in a double auction style market. Two goods are said to be *imperfectly substitutable* if a buyer can use either of them, but prefers one over the other. For example, an electronics manufacturer using a RAM chip can use many suppliers to do this but may be willing to pay a premium for components with a lower failure rate. We give a formal description of a double auction style market mechanism for trading such goods, and define the (classical) equilibrium in such an environment. We present the IS-ZIP agent, which is a generalisation of the ZIP agent for double auctions of Cliff and Bruten. It is able to participate in our double auction environment to make purchases or sales of imperfectly substitutable goods. We demonstrate that, when trading a single good, it is equivalent to Preist and van Tol's modification of the ZIP agent. We describe experiments where a group of IS-ZIP agents with different valuations trade repeatedly, and demonstrate that they rapidly converge to our predicted equilibrium. We conclude by relating our work to that of others, particularly work dealing with multi-attribute negotiation, and discussing extensions.

* Internal Accession Date Only Approved for External Publication
[+] AMEC IV Proceedings, Autumn 2002
[1] Department of Software (Artificial Intelligence), UPC, Barcelona 08034, Spain

Automated Negotiation in Many-to-Many Markets for Imperfectly Substitutable Goods

Chris Preist[1] and Carlos Mérida-Campos[2]

[1] Hewlett Packard Laboratories, Filton Road, Stoke Gifford
Bristol BS12 6QZ, UK
Tel: +44 117 312 8311 Fax: +44 117 312 9250
cwp@hplb.hpl.hp.com

[2] Department Of Software (Artificial Intelligence), UPC
Barcelona 08034, Spain
dmerida@lsi.upc.es

Abstract. In this paper, we present an agent which is able to negotiate the buying and selling of imperfectly sustitutable goods in a double auction style market. Two goods are said to be *imperfectly substitutable* if a buyer can use either of them, but prefers one over the other. For example, an electronics manufacturer using a RAM chip can use many suppliers to do this but may be willing to pay a premium for components with a lower failure rate. We give a formal description of a double-auction style market mechanism for trading such goods, and define the (classical) equilibrium in such an environment. We present the IS-ZIP agent, which is a generalisation of the ZIP agent for double auctions of Cliff and Bruten [3]. It is able to participate in our double auction environment to make purchases or sales of imperfectly substitutable goods. We demonstrate that, when trading a single good, it is equivalent to Preist and van Tol's modification of the ZIP agent [11]. We describe experiments where a group of IS-ZIP agents with different valuations trade repeatedly, and demonstrate that they rapidly converge to our predicted equilibrium. We conclude by relating our work to that of others, particulary work dealing with multi-attribute negotiation, and discussing extensions.

1 Introduction and motivation

Despite the bursting of the dot com bubble, Electronic Commerce has become an increasingly important part of the world economy. More and more transactions, both from business to consumer and between businesses, are taking place online. Despite a downturn in spending on IT in general, companies are continuing to increase investment in e-commerce. For example Computer Weekly/Kew Associates forecast an increase in spending of 7% on e-commerce infrastructure in the UK in 2002.

While simple fixed-cost business transactions can be easily automated using current technology, more complex transactions can involve negotiation to determine price and conditions. Agent technology has been proposed as a means of automating this (e.g. [9]). In this paper, we look at the role of agents in automating negotiations in circumstances where the buyer or seller has several options of what to buy/sell, but prefers some of these options over others. Goods of this nature are referred to as *imperfectly substitutable*. We consider a double-auction style market where many buyers and sellers can meet to trade such goods.

In Section 2, we introduce the problem. We informally define imperfectly substitutable goods, and give a real world example of an internet site which trades them using a double auction. We give a formal description of a double-auction style market mechanism for trading such goods, and define the (classical) equilibrium in such an environment. In section 3, we present the IS-ZIP agent, which is a generalisation of the ZIP agent for double auctions of Cliff and Bruten [3]. It is able to participate in our double auction environment to make purchases or sales of imperfectly substitutable goods. We demonstrate that, when trading a single good, it is equivalent to Preist and van Tol's modification of the ZIP agent ([11],[12]). In section 4, we describe experiments where a group of IS-ZIP agents with different valuations trade repeatedly in our market, and demonstrate that they rapidly converge to our predicted equilibrium. In section 5, we conclude by relating our work to that of others, particulary work dealing with multi-attribute negotiation, and discussing extensions.

2 Markets for Imperfectly Substitutable Goods

When an agent or person is attempting to buy a good or service, they are often faced with many alternate suppliers. Sometimes, these suppliers provide goods which are identical to each other, and as a result the buyer doesn't care which they buy. Usually, however, the suppliers provide slightly different goods. For example, some electronic component suppliers make components which have an exceptionally low failure rate. Manufacturers wishing to have robust products will pay a premium for these, in preference to buying standard quality components. However, if the low-failure rate components are not available for some reason, they could make do temporarily with the standard components. These goods are said to be *imperfectly substitutable*: either good could be used, but the buyer would prefer one over the other. This contrasts with *perfectly substitutable* goods, where the buyer doesn't care which they purchase.

Marketplaces have developed on the web which allow the trading of imperfectly substitutable goods. For example, FastParts (`www.fastparts.com`) provides a double auction where companies can buy or sell electronic components. Sellers can post offers of goods for sale, and buyers can post bids specifying the prices they are willing to pay currently. Both can revise their proposed trade prices based on activity in the market. They can specify features of the goods for sale and the goods required (eg manufacturer, type of packaging, etc). A buyer can select which of the seller offers to move towards or to accept by being more specific about these features, or can remain more general and so potentially trade with one of a wider range of sellers. In this way, buyers and sellers can negotiate about imperfectly substitutable goods.

The FastParts environment has one disadvantage; unless the buyer is willing to risk making multiple purchases accidentally, they can only place one bid. Hence it is not possible to make statements of the form: ‘I am currently willing to pay £1000 for high-quality components, or £700 for similar standard quality components’. Sandholm [14] overcomes this problem in the context of combinatorial auctions through the use of exclusive disjunctive bids. We adopt a similar approach. We now present details of our environment.

2.1 Double Auction for Imperfectly Substitutable Goods

We assume a set of goods, G, labeled $\{g_1 \dots g_n\}$. In each round of the auction, buyers place bids and sellers place offers for some or all of the goods in the set G. We define an *atomic bid* to be a tuple (g_i, p) , representing the fact that the buyer is willing to buy good g_i for at most price p. We define an *exclusive disjunctive bid* or *XOR-bid*, to be a list of atomic bids for different goods, $(g_{i_1}, p_1) \vee \dots \vee (g_{i_m}, p_m)$, representing the fact that the buyer is willing to have at most one of the atomic bids accepted. Similarly, we define an *atomic offer* to be a tuple (g_i, p), representing the fact that the seller is willing to sell good g_i for at least price p, and a *XOR-offer* as a list of atomic offers, at most one of which can be accepted. For the purpose of this paper, we assume that a buyer or seller wishes to trade at most one good at any given time.

Each auction round proceeds as follows:

1. The auction house displays to the buyers and sellers all the bids/offers that remain from the last round (i.e. those which failed to make a trade.) A bid/offer is treated as persisting, unless the trader submits a bid/offer for the same good that improves on the previous one.
2. Buyers submit bids and sellers submit offers to the market. These bids and offers must satisfy a NYSE-style improvement rule: Any atomic bid submitted for a good g must be greater than the highest bid for g persisting from the previous round, and any atomic offer for g must be lower than the lowest offer for g persisting from the previous round. Atomic bids not satisfying these criteria are deleted from XOR-bids containing them.
3. A deal is possible if some XOR-bid contains an atomic bid (g_i, p) and some XOR-offer contains an atomic offer (g_i, q) such that $p \geq q$. If a deal is possible, the auction house identifies all possible deals and selects the one with the largest difference between bid price and offer price. It deletes the corresponding XOR-bid and offer and repeats this process until no more deals are possible. It then informs all participants of the deals made. A deal is made at the price midway between the selected bid and offer.

This marketplace is a simple generalisation of the continuous k-double auction to the multiple good case. Other market mechanisms may be possible, using more sophisticated winner-determination algorithms to clear the market. We do not consider these issues in this paper.

2.2 Equilibria in Markets for Imperfectly Substitutable Goods.

Classical microeconomics provides an approach for determining the equilibrium price in a given market. The quantity of a good that buyers are prepared to purchase at any given price is referred to as the *demand* at that price, and the quantity of a good that sellers are prepared to sell at a given price is the *supply* at that price. According to classical microeconomics, the *equilibrium price* is the price at which supply and demand are equal, and trade should take place at this price when a market stabilizes. Game theoretic work (Such as [13]) has demonstrated that, in general, this is an approximation to reality, and from a theoretical point-of-view requires a large number of traders. This assumption means that the traders act as *price takers*: the action of an individual trader will not effect the equilibrium price. However, experimental work ([15]) has shown that in practice even a small number of traders (5 buyers and 5 sellers) converge rapidly to this equilibrium price, and trade prices do not deviate significantly after stabilization. For that reason, we use this classical definition of equilibrium as the basis of our work.

General Equilibrium Theory of Walras[18] considers a set of markets for many goods, taking into account the substitution effects between them. When the system is in equilibrium, the supply and demand in each market is equal, and no trader wishes to change its supply/demand of any good given the observed equilibrium prices. We now present the definition of a Walrasian equilibrium in the context of our environment.

For each good g_i, let m_i be its associated market. Let M be the set of all such markets. Let S be a set of sellers and B be a set of buyers, all able to participate in any market in M. For a given seller, s, we define r_s to be a function from goods to prices. $r_s(g)$ is the reservation price of seller s for good g. Similarly, for a given buyer b we define v_b to be a function from goods to prices, such that $v_b(g)$ is the valuation of buyer b for good g.

Let $\boldsymbol{A}$: $S \cup B \rightarrow M$ be any total assignment of all buyers and sellers to individual markets in M (i.e. $\boldsymbol{A}$ is a partition of $S \cup B$, and each element in the partition is associated to a market in M). Under such an assignment, we define the supply and demand in market m_i at price p as follows:

$$\begin{aligned} \mathit{supply}(m_i, \boldsymbol{A}, p) &= |\{s \in S | \boldsymbol{A}(s) = m_i \ \& \ r_s(g_i) \leq p\}| \\ \mathit{demand}(m_i, \boldsymbol{A}, p) &= |\{b \in B | \boldsymbol{A}(s) = m_i \ \& \ v_b(g_i) \geq p\}| \end{aligned}$$

In other words, the supply at a given price is the number of sellers assigned to the market under assignment $\boldsymbol{A}$ which are willing to sell at less than or equal to that price, and the demand is the number of buyers assigned to the market willing to buy at greater than or equal to that price. Given this, we can define the equilibrium price as:

$$\mathit{equilibrium\_price\_set}(m, \boldsymbol{A}) = \{p : \mathit{supply}(m, \boldsymbol{A}, p) = \mathit{demand}(m, \boldsymbol{A}, p)\}$$

If the supply function is strictly increasing over price, and the demand function is strictly decreasing over price, then this set will have at most one element, the

equilibrium price. If the supply function is non-decreasing, and the demand function is non-increasing, then this set may be a price range, called the *equilibrium range*. For the purposes of this paper, we assume that there is a unique equilibrium price $equilibrium\_price(m, \boldsymbol{A})$.(Later, we construct experiments which ensure this is the case.) We define the potential profit of a trader t in a market m_i under assignment $\boldsymbol{A}$ as follows:
$profit(t, m_i, \boldsymbol{A}) = v_t(g_i) - equilibrium\_price(m_i, \boldsymbol{A})$ if t is a buyer.
$profit(t, m_i, \boldsymbol{A}) = equilibrium\_price(m_i, \boldsymbol{A}) - r_t(g_i)$ if t is a seller.

If we assume that the traders are price-takers, then they will wish to trade in a market which maximises their profit. Provided this is the market they are assigned to, $\boldsymbol{A}$(t), then the system will be in equilibrium.

Def. 1 *An assignment $\boldsymbol{A}$ is a Walrasian Equilibrium assignment if, for each trader t, there is no market j such that $profit(t, \boldsymbol{A}(t), \boldsymbol{A}) < profit(t, m_j, \boldsymbol{A})$*

In other words, given that the equilibrium price of all markets does not change, no trader wishes to trade in a different market from the one they are allocated. However, in our environment described in section 2.1, traders are not required to make a choice about which good to trade until the actual moment of agreement. For this reason, as we shall see later, we can find situations where there are several equivalent equilibrium assignments which result in identical prices in all markets, and therefore identical profits for all agents. We refer to this equivalence class of assignments as an *abstract equilibrium*, ***Abs***. Formally, it is defined as follows:

Def. 2 *$\boldsymbol{A} \sim \boldsymbol{A'}$ if $\boldsymbol{A}, \boldsymbol{A'}$ are General Equilibria assignments, and $equilibrium\_price(m, \boldsymbol{A}) = equilibrium\_price(m, \boldsymbol{A'})$ for all markets m.
$Abs(\boldsymbol{A}) = \{\boldsymbol{A'} | \boldsymbol{A'} \sim \boldsymbol{A}\}$ is the equivalence class represented by assignment $\boldsymbol{A}$.*

Hence all traders are indifferent with respect to the concrete equilibria within the abstract equilibrium. Even though prices are stable, and the environment is in equilibrium, different numbers of trades can take place in different markets, reflecting trading at different concrete equilibria within the abstract equilibrium.

We have developed a software tool which, given a set of traders, each with valuations/reservation prices for a given set of goods, returns the set of equilibrium assignments. We have used this to determine the theoretical equilibria of the experimental setups described in section 4.

Having defined our market mechanism, and defined what an equilibrium in the market is, we now turn to the problem of constructing agents to act on behalf of buyers and sellers in such a market.

3 The IS-Zip Agent

In the spirit of Cliff & Bruten [3],we have chosen to initially explore the problem of participating in markets for imperfectly substitutable goods by developing an algorithm that is as simple as possible while still able to participate sensibly in trading. Our algorithm is a generalisation of Cliff and Bruten's Zero Intelligence Plus agent. (Or, more precisely, a generalisation of Preist and van Tol's modified ZIP agent to participate in a double auction with order book[11]). The ZIP agent consists of a set of heuristic rules for determining a target shout price (where a buyer shouts a bid and a seller shouts an offer) based on the current state of the market. The actual shout price announced by the agent is determined by applying an adaptation rule to the previous shout and the target shout price. The IS-ZIP agent is designed similarly, but instead of reasoning about a target price, it reasons about a target utility. We now present the IS-ZIP algorithm.

3.1 The IS-Zip Algorithm

The utility of a potential trade of good g_i at a price p is defined as:

$tradeUtility(g_i, p) = v_t(g_i) - p$ if t is a buyer.

$tradeUtility(g_i, p) = p - r_t(g_i)$ if t is a seller.

The goal of the agent is to maximise its trading utility, by trading at most one good. In the experiments described in section 4 it repeats this, being reinitialised with permission to trade at the start of each new market 'day'.

The following pseudocode algorithm is used to set the target bid value for the agent, assuming the agent is a buyer:

```
IF(TradingNow & IMustTrade)
    NewTarget := Maximum( BestSellerDealUtility +  δ ,
                          BestBuyerShoutUtility - δ);
ELSEIF(¬ TradingNow & IMustTrade)
    NewTarget := BestBuyerShoutUtility - δ;
ELSEIF (TradingNow & ¬ IMustTrade)
    NewTarget := Maximum( OldTarget, BestSellerDealUtility +δ);
ELSE NewTarget := OldTarget;
```

Where:

- `TradingNow`: Boolean variable set to TRUE when trades were made in the last market round.
- `IMustTrade`: Boolean variable set to TRUE when the agent wishes to make a purchase.
- `OldTarget`: Real variable set to the value of the algorithms target in the last market round.
- `BestSellerDealUtility`: The maximum utility to the agent of the offers made by sellers which were accepted last round.
- `BestBuyerShoutUtility`: The maximum utility to the agent of the highest bid made last round for each good g_i.
- δ: An arbitrary (possibly randomly determined) small amount.

Formally, we can define `BestSellerDealUtility` and `BestBuyerShoutUtility` as follows: Let B be the set of atomic bids made last round, O the set of atomic offers, and A the set of atomic bids and offers which were accepted. (Disjunctive bids and offers are split into their component atomic bids and offers for the purpose of this definition.)

$$\texttt{BestSellerDealUtility}= max\{tradeUtility(g_i, p)|(g_i, p) \in O \cap A\}$$
$$\texttt{BestBuyerShoutUtility}= max\{tradeUtility(g_i, p)|(g_i, p) \in B \; \& \; \& \; [(g_i, q) \in B \rightarrow q \leq p]\}$$

The algorithm for the seller is identical to this, except that `BestSellerDealUtility` is replaced with `BestBuyerDealUtility`, and `BestBuyerShoutUtility` is replaced with `BestSellerShoutUtility`, defined in the obvious way.

Given the target value, the agent calculates the actual utility value of the bids to place in the same way as Cliff's ZIP agent, by using the Widrow-Hoff rule with momentum:

Let $u(time)$ be the target utility for the agent at a given time $time$, and $q(time-1)$ be the actual utility used at $time-1$ to calculate the bids/offers placed by the agent. The actual utility used at $time$ is given by $q(time) = q(time-1)+\Gamma(time)$, where $\Gamma(time) = \gamma\Gamma(time-1)+(1-\gamma)\Delta(time-1)$. γ is the *momentum coefficient* for the agent, and $\Delta(time) = \beta(u(time)-q(time-1))$. β is the *learning rate* for the agent.

¿From the utility value, $q(time)$, returned by the adaptation rule, the agent submits a disjunctive bid defined as follows:

$$\bigvee_{g_i \in G} (g_i, v_t(g_i) - q(time)) \qquad \textit{if t is a buyer}$$
$$\bigvee_{g_i \in G} (g_i, r_t(g_i) + q(time)) \qquad \textit{if t is a seller}$$

3.2 The (PS)ZIP Agent as a Special Case of the IS-ZIP Agent

If we deploy the IS-ZIP agent in an environment where only a single good is traded, then we can demonstrate that the set of heuristic rules are equivalent to those used in Preist and van Tol's modification of the ZIP agent for use in a double auction with order book (The PS-Agent)[11][12].

Theorem 1 *The IS-ZIP agent applied to a market with a single good is equivalent to the PS-Agent.*

Proof : We present the proof for the buyer case. The seller case follows through in a similar fashion.

Consider firstly the definition of $BestSellerDealUtility$:

$$\begin{aligned} BestSellerDealUtility &= max\{tradeUtility(g_i, p) \,|\, (g_i, p) \in O \cap A\} \\ &= max\{v_t(g) - p \,|\, p \in O \cap A\} \\ &\quad \textit{in the case where there is a single good.} \\ &= v_t(g) - min\{p \,|\, p \in O \cap A\} \end{aligned}$$

$$= v_t(g) - min\{p \mid p \in O\}$$
$$given\ that\ A\ is\ non-empty.\ (Because\ the\ lowest\ price\ offer\ will\ certainly\ be\ in\ A\ in\ this\ case.)$$

Similarly, BestBuyerShoutUtility simplifies as follows:

$$\begin{aligned} BestBuyerShoutUtility &= max\{tradeUtility(g_i, p) \mid \\ &\quad \mid (g_i, p) \in B \ \wedge\ [(g_i, q) \in B \rightarrow\ q \leq p]\} \\ &= max\{v_t(g) - p \mid p \in B \ \wedge\ [q \in B \rightarrow\ q \leq p]\} \\ &= v_t(g) - min\{max\ p \in B\} \\ &= v_t(g) - max\{p \in B\} \end{aligned}$$

Now consider the first assignment of NewTarget in the algorithm specification:

```
IF(TradingNow & IMustTrade)
    NewTarget := Maximum( BestSellerDealUtility +  δ ,
                          BestBuyerShoutUtility - δ);
```

In this case, the definition simplifies as follows:

$$\begin{aligned} &Maximum\ (\ BestSellerDealUtility + \delta\,,\ BestBuyerShoutUtility - \delta) \\ &\qquad = max(v_t(g) - min\{p|p \in O\} + \delta\,,\ v_t(g) - max\{p \in B\} - \delta) \\ &\qquad = v_t(g) - min\{p|p \in O\} + \delta \\ &\qquad\quad because\ min\{p|p \in O\} \leq max\{p \in B\} \\ &\qquad\quad if\ trades\ are\ taking\ place. \end{aligned}$$

Inserting this, together with the previous simplifications, into the algorithm definition, we get:

```
IF(TradingNow & IMustTrade)
    NewTarget := v_t(g) - min{p | p ∈ O} + δ;
ELSEIF(¬ TradingNow & IMustTrade)
    NewTarget := v_t(g) - max{p | p ∈ B} - δ;
ELSEIF (TradingNow & ¬ IMustTrade)
    NewTarget := Maximum(OldTarget, v_t(g) - min{p | p ∈ O}  + δ);
ELSE NewTarget := OldTarget;
```

These define the target utilities. The definition of PS-Agent in [12] is given in terms of the target shout price. The IS-ZIP agent shouts a bid of $v_t(g) - u$, so we can fold this in to the above definition:

```
IF(TradingNow & IMustTrade)
    NewTargetBid := min{p | p ∈ O} - δ;
```

```
ELSEIF(¬ TradingNow & IMustTrade)
    NewTargetBid := max{p | p ∈ B} + δ;
ELSEIF (TradingNow & ¬ IMustTrade)
    NewTargetBid := Minimum(OldTarget, min{p | p ∈ O}  - δ);
ELSE NewTarget := OldTarget;
```

This is equivalent to the definition of PS-Agent given in [12]. □

4 Experimental Analysis

We now describe experiments which deomnstrate empirically that a community of IS-ZIP agents trading in a variety of supply/demand conditions result in convergence of trade towards the predicted equilibrium patterns described in section 2.2.

4.1 Experimental Procedures

Our experimental setup is based on that used by Smith [15] with human subjects, and later applied by Cliff and Bruten [3] to ZIP agents. We assign valuations and reservation prices to a set of buyers and a set of sellers, who repeatedly trade in a double-auction environment as described in section 2.1. Each experiment is divided into 500 'days'. At the beginning of each day, and agent is assigned the task of buying/selling one item. Each day is divided into 'rounds', where buyers place bids and sellers place offers, with trade being determined by the double auction mechanism. Each day consists of 100 such rounds. If an agent trades during a day, it receives a utility credit and can make no further trade until the next day. If an agent makes no trade in a day, it receives no credit.

The double auction trades goods of three different qualities : L(ow), M(edium), H(igh). Agents have different valuations/reservation prices for goods of different quality, hence they are imperfectly substitutable. Different experimental setups, described below, assign different such values to agents. For the purposes of the experiments presented in this paper, all agents have identical learning parameters ($\alpha = 0.4, \Gamma = 0.1$) as well as identical initial target utilities (40 for sellers, and 20 for buyers).

To measure the convergence of the markets towards the equilibrium, we use a generalisation of *Smith's alpha* [15]. Smith's alpha is defined for a standard double auction as the standard deviation of the observed trade prices from the theoretical equilibrium price, expressed as a percentage:

$$\alpha = \sqrt{\sum_{i=1}^{n} \frac{\left(\frac{P_i - P_O}{P_O}\right)^2}{n}}$$

Where P_i is the price at wich deal i has been made. P_O is the equilibrium price computed 'a priori', and n is the total number of deals that have been made. The smaller this value is, the closer trade is taking place to equilibrium.

We define the *unified alpha* to be a generalization of this measure to compute the standard deviation for a set of related markets:

$$\alpha = \sqrt{\sum_{j=market_0}^{market_m} \left(\sum_{i=1}^{n_j} \frac{\left(\frac{P_{ij}-P_{Oj}}{P_{Oj}}\right)^2}{n} \right)} \qquad \text{Where } n = \sum_{j=0}^{m} n_j$$

Where n_j is the number of trades in market j, P_{ij} is the price of trade i in market j and P_{Oj} is the equilibrium price of market j. This can be shown to be equivalent to:

$$\alpha = \sqrt{\sum_{j=market_0}^{market_m} \left(\frac{n_j}{n} \left(\alpha_j^2\right) \right)}$$

Where α_j is Smith's alpha for market j.

4.2 Experimental Set-Ups

Each experimental set-up consists of a different assignment of valuations to a set of agents. In most experiments, we use 20 buyers and 20 sellers, though in some we use only 10 of one category. In table 2, we present the valuations for the basic experimental setup. In table 1, we present other setups as modifications of this basic setup.

Table 1. Market characteristics in the different experimental set-ups

Experiment	**Market Characteristics**
Basic set up (*BS*)	Agents given valuations in table 2.
Reduced demand (*RD*)	Sellers have same valuations as BS. Buyers reduce valuations by 10 for all goods.
Increased supply (*IS*)	Buyers have same valuations as BS. Sellers reduce valuations by 10 for all goods.
Increased demand for High quality (*IDH*)	Sellers have same valuations as BS. Buyers have same valuations for L and M goods, and increase valuation of H goods by 10.
Reduced number of sellers (*RS*)	20 buyers as BS. 10 sellers, valuations given in table 3.
Reduced number of buyers (*RB*)	20 sellers as BS. 10 buyers, valuations given in table 3.

We have used the algorithm described in section 2.2 to determine the abstract equilibrum for each experimental setup. Table 5 gives the abstract equilibrium for the basic setup. Table 5 presents the average price of all deals in each market over a 500 day experiment, and the average number of trades in each market.

Table 2. Reservation prices for the basic set up.

Buyer	Qty : L	Qty : M	Qty : H	Seller	Qty : L	Qty : M	Qty : H
0	10	30	40	0	10	30	40
1	10	30	40	1	10	30	40
2	45	55	60	2	45	55	65
3	45	55	60	3	45	55	65
4	50	60	65	4	65	65	70
5	50	60	65	5	65	65	70
6	55	65	70	6	65	75	95
7	55	65	70	7	65	75	95
8	60	70	75	8	70	80	85
9	60	70	75	9	70	80	85
10	65	75	80	10	75	85	90
11	65	75	80	11	75	85	90
12	70	80	85	12	80	80	80
13	70	80	85	13	80	80	80
14	75	85	90	14	80	90	95
15	75	85	90	15	80	90	95
16	80	90	95	16	85	95	100
17	80	90	95	17	85	95	100
18	85	90	100	18	100	105	110
19	85	95	100	19	100	105	110

Table 3. Reservation prices for the Buyers in RB and sellers in RS.

Buyer	Qty : L	Qty : M	Qty : H	Seller	Qty : L	Qty : M	Qty : H
0	10	30	40	0	45	65	75
1	10	30	40	1	50	60	70
2	45	55	60	2	65	65	70
3	50	60	65	3	60	70	90
4	55	65	70	4	75	85	90
5	60	70	75	5	80	80	80
6	65	75	80	6	80	90	95
7	70	80	85	7	85	95	100
8	75	85	90	8	100	105	110
9	80	90	95	9	100	105	110

Table 4. Results of 500 days 100 rounds negotiation with basic set-up.

	Quality : L	Quality : M	Quality : H
Deals	2014	1465	1521
Avg. price	64.725	74.768	79.935
Avg. deals per day	4.028	2.93	3.042

Table 5. Theoretical equilibrium for the basic set-up.

Equilibrium issue	Value
Number of deals per day	10
Equilibrium price for Qty.L market	65
Equilibrium price for Qty.M market	75
Equilibrium price for Qty.H market	80
Abstract equilibrium pattern	2 deals with quality 10 2 deals with quality 30 4 deals with quality 10 or 20 2 deals with quality 20 or 30

Fig. 1. Daily profit of Buyer 15 (20 days) and Unified alpha (500 days) using the BS

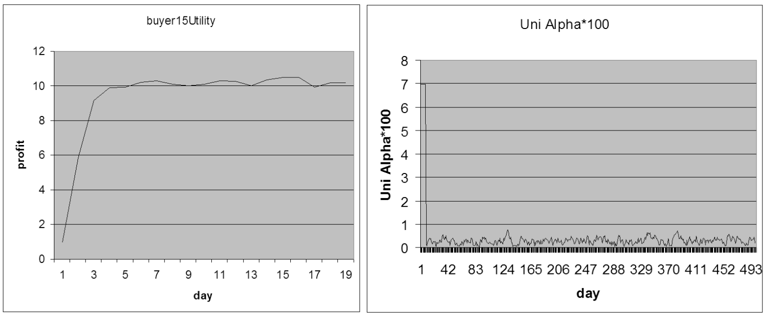

Fig. 2. Unified alpha for each configuration for the first 10 days.

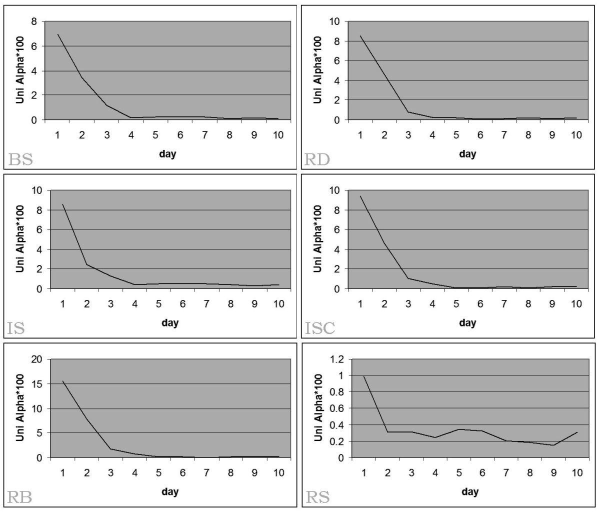

Table 6. Equilibrium test:Predicted prices vs. Average prices.

Set Up	**Pred. price(L, M, H)**	**Avg. price(L, M, H)**
BS	(65, 75, 80)	(64.7, 74.7, 79.9)
IS	(60, 70, 75)	(59.6, 69.6, 74.6)
RD	(60,70,75)	(59.9, 69.9, 74.95)
IDH	(60,70,85)	(59.9, 70.0, 84.9)
RB	(55,65,70)	(55.0,65.0,70.0)
RS	(75,85,90)	(74.5,84.5,89.6)

As we can see, the average price is close to the theoretical equilibrium price, suggesting convergence. This is confirmed by figure 1 (right), which shows how the unified alpha rapidly reduces to be below 0.5, and remains there. More detail of this convergence is given in the first graph in figure 2. Figure 1 (left) shows the daily profit of one of the agents (Buyer 15) during the first ten days. It shows its profit increasing towards that predicted by the equilibrium. We have also carried out an analysis of the *trading pattern* each day - the distribution of trades between the three markets. We have determined that every trading pattern in the 500 days was within the abstract equilibrium predicted, with the most common being 4 trades of L, and 3 each of M and H. This occured 98 times over the 500 days. As a result of this, we can conclude that the IS-ZIP agents in the basic setup do successfully converge to the predicted equilibrium.

We have carried out similar analysis of the other 5 setups, which we present in less detail for reasons of space. Table 6 shows the predicted equilibirum prices for each setup, together with the observed average over 500 days. Figure 2 presents the unified alpha for each experimental setup. This demonstrates that in all of our experiments the market converges to equilibrium over a period of 4-5 days.

These results strongly suggest that communities of IS-ZIP agents rapidly converge to the equilibrium when trading imperfectly substitutable goods.

5 Discussion/Conclusions

In this paper, we have presented a simple adaptive agent which is able to negotiate to buy or sell one of a set of imperfectly substitutable goods. We have provided initial results that suggest that it behaves reasonably, and converges rapidly to the equilibrium predicted by our analysis in section 2.2. As far as we are aware, this is the first reported agent able to exploit the economic properties of a double auction for imperfectly substitutable goods. The agent we developed is based on the ZIP agent of Cliff and Bruten [3].This is a non-optimal strategy in the market, but nonetheless has been shown to outperform humans in experiments [4]. Hence, we believe that a generalisation of ZIP is a valuable first step in the study of agents in markets containing imperfectly substitutable goods. However, we do not view this as the definitive solution. Other agent designs have been developed for deployment in double auctions trading a single commodity good. He et. al. [8] use a set of fuzzy heuristics to determine bid and offer prices based on past history. Gjerstad and Dickhaut [7] use function fitting over prior history, together with utility analysis, to determine an effective bid/offer to make. Both of these have been shown to outperform ZIP agents in certain environments ([8], [4]). Tesauro and Das [16] have modified the Gjerstad-Dickhaut algorithm, resulting in further performance improvements over the other strategies. Tesauro and Bredin [17] use a dynamic programming which outperforms this. Park et. al. [10] propose using a stochastic-based algorithm. Walsh et. al. [19] analyse strategic interactions in the choice of strategies. Any of these strategies may provide a basis for developing more sophisticated agents able to handle imperfectly substitutable goods. We are currently exploring extensions of [8] and [16].

Cheng and Wellman ([2]) developed WALRAS, an agent-based distributed algorithm which is able to find the general equilibrium of an economy of many goods, where traders have a continuous utility function. The approach it uses is based on 'tattonement' of Walras ([18])where markets for each good announce an interim clearing price, and agents adjust their supply/demand of the goods based on these price signals. This process repeats until the clearing prices reach equilibrium. Cheng and Wellman have proved that their algorithm converges when preferences on goods are strictly convex (i.e. when any pair of goods is more valuable than either of the individual goods), and when goods are grossly substitutable (i.e. An increase in the price of one good will not lead to a reduction in the demand for another.)

Our work differs from the WALRAS algorithm in that trading can take place prior to the equilibrium being established, and that the equilibrium is found through observation of point trades rather than the announcement of provisional clearing prices. Our assumption that traders wish to trade exactly one good violates the assumption that preferences are strictly convex, and also that the utility function of traders is continuous. Hence the convergence result of Cheng and Wellman cannot be applied directly to our system. We hope to explore the relationship between the two systems more fully in the future.

The problem of negotiating the trade of an imperfectly substitutable good is related, though not equivalent to, the problem of negotiating the trade of multi-parameter goods. A multi-parameter good has other parameters beyond price which must be agreed between trading parties, such as delivery time, colour, quality, etc. The experiments presented above could be viewed as experiments in (very limited) multi-attribute negotiation where the agents negotiate price, which is a (near-) continuous parameter, and quality, which is a discrete parameter with only 3 values. More generally, we need to make two assumptions if our agents are to be deployed in a multi-parameter environment. Firstly, we must assume that all agents announce disjunctive bids containing atomic bids for all possible sets of parameter assignments. Secondly, we must assume that there is a sufficiently large agent population with sufficiently varied valuations over the space of possible parameter assignments. Together these two assumptions will result in the (imaginary) market for each set of parameter assignments having a unique equilibrium price, or at least a narrow range of equilibrium prices. This allows our algorithm to have enough information about the current state of the market to adjust its expectations appropriately. Clearly, if the number of possible parameter assignments is large, this will become impractical, and if a non-price parameter is continuous, it is impossible.

This is a very different, but complementary, view on multi-parameter negotiation from the work carried out by the agent community up to now. This work has tended to focus on one-to-one negotiation. Zeng and Sycara [20] present an approach based on using learning to build a model of the trading partner's reservation price. Faratin et. al. [5] use heuristics to make appropriate concessions based on the passage of time, use of resources, and other factors. They augment this approach with fuzzy reasoning [6], to identify the set of parameters which are most likely to be acceptable to the trading partner based on past experience. Barbuceanu and Lo [1] use Multi-Attribute Utility Theory to build a structured model of the preferences an agent has, to allow it to make concessions in an appropriate order. All these approaches differ from our work, in that they focus primarily on the one-to-one environment. They could be deployed in an environment with many buyers and sellers, but will not be able to exploit the economic dynamics created by seller-to-seller and buyer-to-buyer competition. However, they have the advantage over our approach that they are able to operate in environments where there is a far lower amount of information available - i.e. there is only a single proposal from the trading counter-party, detailing its current proposed price for only one set of parameter values.

We believe that the way forward for multi-parameter negotiation should combine the best of these two approaches - it should exploit economic dynamics to determine the range of possible prices for given parameter values, but also use heuristic techniques to concede appropriately within these ranges, and fuzzy reasoning like that of [6] to determine which parameter values to propose. We believe that the negotiation mechanism should allow disjunctive bids and offers, to enable traders to simultaneously propose prices for several sets of parameter values, but should not force traders to make bids/offers for all possible parameter values, as our algorithm does.

References

1. Barbuceanu,M. and Lo,W. A multi-attribute utility theoretic negotiation architecture for electronic commerce" Proceedings of 4th Int. Conf. on Autonomous Agents, Barcelona, Spain, 239-247
2. Cheng, J. and Wellman, M. The WALRAS Algorithm: A Convergent Distributed Implementation of General Equilibrium Outcomes. Computational Economics 12: 1-24, 1998.
3. Cliff, D. and Bruten, J. Less than Human: Simple Adaptive Trading Agents for CDA Markets. Proceedings of the 1998 Symposium on Computation in Economics, Finance and Engineering: Economic Systems, 1998.

4. Das,R., Hanson, J.E., Kephart,J.O., and Tesauro, G. Agent-Human Interactions in the Continuous Double Auction. Proceedings of the International Joint Conference on Artificial Intelligence, Seattle, 2001
5. Faratin,P., Sierra,C., and Jennings,N.R. Negotiation Decision Functions for Autonomous Agents. Int. Journal of Robotics and Autonomous Systems, 24 (3-4) 159-182, 1998
6. Faratin,P., Sierra,C., and Jennings,N.R. Using similarity criteria to make negotiation trade-offs. Proc. 4th Int. Conf. on Multi-Agent Systems (ICMAS-2000), Boston, USA, 119-126, 2000.
7. Gjerstad, S. and Dickhaut,J. Price Formation in Double Auctions. Games and Economic Behaviour, 22:1-29, 1998.
8. He,M., Leung, H. and Jennings, N.R. A Fuzzy Logic Based Bidding Strategy for Autonomous Agents in Continuous Double Auctions. To appear in IEEE Trans on Knowledge and Data Engineering.
9. Jennings,N.R., Faratin,P., Norman,T.J., O'Brien,P., Wiegand,M.E., Voudouris,J., Alty,J. L., Miah,T., and Mamdani, E. H., ADEPT: Managing Business Processes using Intelligent Agents, in: Proc. BCS Expert Systems 96 Conference (ISIP Track), Cambridge, UK pp5-23, 1996
10. Park,S., Durfee,E. and Birmingham,W. An Adaptive Agent Bidding Strategy based on Stochastic Modelling. Proc. 3rd International Conf on Autonomous Agents (AA'99), ACM, May 1999 pp 147-153
11. Preist, C. and van Tol, M. Adaptative Agents in a Persistent Shout Double Auction. Proceedings of the 1st International Conference on the Internet Computing and Economics, 1998.
12. Preist, C. Commodity Trading using an Agent-Based Iterated Double Auction. Proc. 3rd International Conf on Autonomous Agents (AA'99), ACM, May 1999 pp 131 - 138
13. Rustichini,A., Satterthwaite, M.A. and Williams, S.R. Convergence to Efficiency in a Simple Market with Incomplete Information. Econometrica Vol 62(5) 1041-1063, 1994.
14. Sandholm, T. eMediator: A Next Generation Electronic Commerce Server. Proc. International Conference on Autonomous Agents (AGENTS) , Barcelona, Spain, 2000.
15. Smith, V. An Experimental Study of Competitive Market Behaviour. Journal of Political Economics, 70, 111-137, 1962.
16. Tesauro, G. and Das, R. High-Performance Bidding Agents for the Continuous Double Auction. Proceedings of IJCAI workshop on Economic Agents, Models and Mechanisms, 2001.
17. Tesauro, G. and Bredin, J. Strategic Sequential Bidding in Auctions using Dynamic Programming. Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems. ACM Press, 2002.
18. Walras, L. Elements of Pure Economics, 1874. Tr. W. Jaffe. Allen and Unwin, 1954.
19. Walsh, W.E., Das, R., Tesauro, G. and Kephart, J.O. Analysing Complex Strategic Interactions in Multi-Agent Games. In Proceedings of AAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT-02), July 2002.
20. Zeng,D. and Sycara,K. Bayesian learning in negotiation. In Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, pages 99–104, Menlo Park,CA, March 1996. AAAI Press