A Generic Framework for a Combined Agent-based Market and Production Model

Bas Straatman · Danielle J. Marceau ·
Roger White

Accepted: 31 August 2012 / Published online: 18 October 2012
© Springer Science+Business Media New York 2012

Abstract Agent-based market models are in general based on a-priori defined supply and demand schemes. Likewise, production models assume that prices are known a-priori. In reality prices depend on variable demands and supplies, while demand and supply depend on variable prices, and these two processes are interconnected. This paper describes a model that for the first time simulates a combined agent-based double auction market and production model. The model is built around von Neumann technology matrices (von Neumann, *Rev Econ Stud* 13(1):1–9, 1946) which provide the links between products. Agents possess one or more technologies to produce products from other products. They trade in order to acquire the inputs and sell in order to generate revenue, and the price is determined by a process of negotiation between buyers and sellers. The algorithm of negotiation is based on Cliff's Zero Intelligence Plus approach (Cliff, *Minimal-intelligence agents for bargaining behaviors in market-based environments*. Technical report, School of Cognitive and Computing Sciences, University of Sussex, 1997), but instead of a single commodity with fixed limit prices the agents change their limit prices for multiple products based on the simulated economic situation. The combination of production and market provides a simple but complete bottom-up model framework for microeconomics. As the results show, the model employs a price mechanism that results in an appropriate allocation of resources without a central command.

B. Straatman (✉)
Energy and Environmental Systems Group, ISEEE, University of Calgary,
Calgary, AB T2N 1N4, Canada
e-mail: bastraat@ucalgary.ca

D. J. Marceau
Department of Geomatics Engineering, University of Calgary, Calgary, AB T2N 1N4, Canada

R. White
Department of Geography, Memorial University of Newfoundland, St. John's, NL A1B 3X9, Canada

Keywords Agent-based computational economics · Von Neumann technology matrices · Markets · Price mechanism · Invisible hand

1 Introduction

Price dynamics in many markets (stocks, resources, housing to name a few) get much attention in the news, every day. As a result this is a very active research area. Recently agent-based approaches to price dynamics have become increasingly popular as a tool to study the behaviour that underlies the observed market and price dynamics. Of the many different forms of markets or auctions, this paper focuses on the Continuous Double Auction (CDA). The CDA allows a group of sellers and a group of buyers to establish prices by means of negotiation rounds where both buyer and seller can announce a bid or an offer at any time. Agent-based approaches to price dynamics of CDAs have led to computerized or automated zero intelligence (ZI) traders instead of human traders.

It has been shown that slightly more advanced versions of these automated traders, zero intelligence plus (ZIP) traders, can outperform human traders in a CDA. However, the existing work on ZIP traders also has its limitations. The agent traders of [Gode and Sunder \(1993\)](#) and [Cliff \(1997\)](#) are limited to a single commodity and make use of fixed limit prices. Obviously, reality is more complicated. The price of a commodity depends not only on the demand and supply of the commodity itself, but also on changes with the price of complementary products. Furthermore, the price of a commodity ultimately depends on the cost of production.

In neoclassical microeconomic theory, price and production (output) are treated as a single, integrated problem. However, current agent-based approaches to price dynamics abstract the problem of price determination from the economic system in which the price mechanism is embedded by assuming the supply of tradable goods. Consequently these models are of limited relevance, applying only to very short term situations or to cases like equities markets where supply is essentially independent of price. Here we return to a unified approach by presenting an agent-based model that for the first time combines a production model with a double auction market model. The production model is built around von Neumann technology matrices ([von Neumann 1946](#)); the market model of demand, supply and negotiations is an advanced version of the zero intelligence plus (ZIP) traders of [Cliff \(1997\)](#). As a result, prices determine production and production determines pricing. The combination of an agent-based production and market system provides a simple but complete bottom-up model framework for an agent-based approach to microeconomics. It thus implies that the general equilibrium predicted by the neoclassical approach is robust to the relaxation of several of the implicit assumptions underlying that approach. For example, as we will show, an equilibrium exists even under imperfect competition conditions, when a very limited number of heterogeneous firms with shifting production plans with bounded rationality get together at a CDA.

And not only does an equilibrium exist. This automated decentralized negotiation process generates a market-based control of the production process. Here we present an autonomous decentralized agent-based model that consists of a production model

and a double auction market model. Instead of using a Walrasian auctioneer that oversees the market and sets appropriate prices in order to balance supply and demand as presented in [Straatman et al. \(2008\)](#), here the agents with their negotiation strategy are fully autonomous as there is no central command, and thus the agents collectively generate the prices for the set of products.

What is new about the coupling of the market and production models is that:

1. negotiating agents change their limit prices based on what is happening in the production system;
2. the model handles multiple commodities, so price changes not only follow changes in demand and supply, but also reflect changes in the pricing of complementary goods;
3. the prices resulting from the rounds of negotiation depend on supply and demand, but in turn they also direct supply and demand through their complete control of the production system.

Furthermore, the coupling of the market and production models gives a more natural implementation of the origin and replenishment of goods. The replenishment schemes such as the continuous and the periodic replenishment in [Cliff and Preist \(2001\)](#) and [Chaggar et al. \(2008\)](#) are convenient for experimentation, but are rather random in their nature. The agents (human in the case of [Smith \(1962\)](#), ZIP traders in the case of [Chaggar et al. \(2008\)](#) and [Cliff and Preist \(2001\)](#)) have their entitlements (units to sell or money to make purchases) reset regularly, or they receive units or money at random times. But where do these resources come from and when do they appear? Obviously goods have to be produced and revenue has to be generated, and the integration of a price formation model with a production model is the only natural way to replenish resources in a continuous double auction market process.

Therefore, the objective of this paper is to present a new integrated stable model of market and production that: (a) can handle multiple products in recognition of the fact that prices of complementary goods are interdependent, (b) incorporates the strong two-way relation between price and supply and demand, (c) adds to the state of the art of agent-based CDA modelling by introducing a natural implementation of the origin and replenishment of goods and a natural implementation of a market-based control that is both automated and decentralized. Finally, the stability of the integrated model is guaranteed by designing the market in such a way that it provides the production system with timely and clear signals, so that the production of all required commodities occurs on a continuing basis. Such an agent-based version of neo-classical micro-economic theory will be very useful in exploring the underlying drivers of the complex economic system that we find ourselves in.

2 Overview of Existing Work on Agent-based Modelling of Price Dynamics in the CDA

The Marshallian cross is the familiar diagram that shows the relation between supply and demand ([Canterbery 2001](#)), with price and quantity sold resulting from the equilibrium of supply and demand. As the price increases, supply goes up while demand goes down; as the price decreases, supply goes down and demand goes up. It is assumed that

there are many buyers and sellers at each price and thus that the curves are continuous. If at low price demand is higher than supply, and at high price supply is higher than demand, these two curves intersect at an equilibrium price and quantity: the price at which demand equals supply.

It is assumed that the market mechanism is capable of locating this equilibrium price at which the market is in balance. Walras introduced the idea of an auctioneer, someone who makes an inventory of calls to buy and sell, provides information to buyers and sellers based upon which a new round of calls is placed; this is repeated until supply matches demand and the actual transactions are allowed to proceed (Canterbery 2001).

Although some transactions are regulated by auctioneers, many more seem to have a more spontaneous nature. Buyers and sellers have a feel for the market, and place their calls based on what they sense is going on. The market organisation here is closely mimicked by the continuous double auction procedure. This particular market design has received much attention for two reasons: this design resembles major financial markets such as the London and New York stock exchanges, and secondly this design is capable of reaching competitive equilibrium faster and more efficiently than other types of auctions (Davis and Holt 1993).

Smith (1962) executed a number of double auction trading experiments in his classroom in the 1960s. Students were handed out cards indicating if they were acting as a buyer or a seller and indicating a limit price. Limit prices were private information not known to the other players and buyers were not allowed to buy above their limit price, while sellers could not sell below their limit price. A transaction at the limit is better than no transaction at all, but obviously the bigger the difference between the limit price and the actual transaction price, the better. In the case where an offer and a bid cross—i.e. the offer price is lower than the bid price—the traders agree on a price and a transaction takes place. Smith's study showed that even with a very limited number of participants (usually the group of buyers and sellers was made of around 20 to 40 students) without perfect knowledge of supply and demand, the double auction mechanism quickly led to convergence to the theoretical equilibrium price that follows from the distributions of limit prices for buyers and sellers.

Gode and Sunder (1993) investigated whether the capacity of human traders to approach equilibrium depends on human intelligence. According to their work, Smith's results can be emulated by artificial traders with Zero Intelligence (ZI). Artificial agents placed random bids, and if the bids were constrained by limit prices as was the case in Smith's experiments, the transaction prices quickly converged to equilibrium. Gode and Sunder therefore concluded that intelligence is not required; the convergence to equilibrium is merely a consequence of the structure provided by the applied double auction market mechanism.

In a subsequent study, Cliff (1997) showed that there is something fundamentally wrong with Gode and Sunder's experimental configuration and that zero intelligence is not enough. The rapid convergence of transaction prices by the constrained zero intelligence agents is a consequence of the probability distributions of the bids and offers. Coincidentally the level at which transaction prices converge is approximately the same as the theoretical equilibrium price that follows from the intersection of supply and demand. Cliff showed that with different demand and supply slopes, and

thus different equilibrium prices, transaction prices of trade between zero intelligence agents does converge, but not at the price level expected based on the supply and demand curves. Thus, Cliff creates the Zero Intelligence Plus (ZIP) agent.

Instead of random bids and offers constrained by a limit price, Cliff proposes bids and offers that depend on the limit price and an agent's personal profit margin. It is the profit margin that has to be adapted up or down based upon what happens in the market (reproduced from [Cliff \(1997\)](#)):

- For sellers:
 - if the last shout was accepted at price q ,
 - then:
 - any seller s_i for which the limit price $p_i \leq q$ should raise its profit margin
 - if the last shout was a bid, then:
 - any active seller s_i for which $p_i \geq q$ should lower its margin,
 - else:
 - if the last shout was an offer, then:
 - any active seller s_i for which $p_i \geq q$ should lower its margin.
- For buyers:
 - if the last shout was accepted at price q
 - then:
 - any buyer b_i for which $p_i \geq q$ should raise its profit margin
 - if the last shout was an offer, then:
 - any active buyer b_i for which $p_i \leq q$ should lower its margin,
 - else:
 - if the last shout was an bid, then:
 - any active buyer b_i for which $p_i \leq q$ should lower its margin.

When agents apply these adaptations to their profit margins as the negotiation rounds take place, the transaction prices quickly converge to the theoretical equilibrium price, including the cases for which the method of [Gode and Sunder \(1993\)](#) results in mismatches. Cliff developed a very elegant model for market clearing. Whether or not it resembles the process of human traders he leaves for others to explore.

Literature provides a variety of agent-based simulations of CDA's. First of all there is an abundance of literature discussing the CDA. For an overview see for example [Friedman and Rust \(1993\)](#) and [Davis and Holt \(1993\)](#). All agree that the CDA is effective in reaching closing prices where supply equals demand, but what is not so clear is how this equilibrium is reached and what affects the dynamics of reaching this equilibrium. For Cliff this is one of the reasons to apply software 'agents' or 'traders' ([Cliff 1997](#)) to the CDA. One can distinguish several other reasons. [LeBaron et al. \(1999\)](#) and [LeBaron \(2001\)](#) for example focus on applications in financial markets. Others focus on automated internet commerce and trading strategies ([Chen 2000](#); [Das 2003](#); [Cliff 2006](#); [Ma and Leung 2008](#)). Others again address specific questions about observed dynamics in trading ([Cliff and Preist 2001](#); [Chiarella and Iori 2002](#)). But none offers an overarching general framework to study market-based control (MBC) of resource allocation. According to [Cliff \(1997, p. 2\)](#), "no current MBC systems are both automatic and distributed".

Table 1 Input and output matrices representing technologies and the commodities involved

Commodities	Input matrix					Output matrix				
	t_1	t_2	t_3	t_4	t_5	t_1	t_2	t_3	t_4	t_5
c_1	1	1	1	1	0	0	0	0	0	2
c_2	0	0	0	0	0	0	0	0	0	0
c_3	0	1	0	0	0	1	0	0	0	0
c_4	0	0	1	0	0	0	2	0	0	0
c_5	0	0	0	1	0	0	0	2	0	0
c_6	0	0	0	0	1	0	0	0	$\frac{7}{5}$	0

Real markets integrate many components with many different actors and they are dynamic. The demand and supply curves of butter will be different tomorrow. And the price of butter depends not only on the demand and supply of butter, but changes with the prices of bread and margarine, complimentary and substitute goods (Kauffman 2000, p. 212). Furthermore, the price of butter depends on the price of milk, an input. And where and when are the goods produced? Obviously the goods have to come from somewhere, but random replenishment schemes that introduce random numbers of a good to be traded at the beginning of the day or at random times as in Chaggar et al. (2008) are ad hoc. The interconnectivity that exists between products, as well as the production system of these products need to be taken into account. A combined production model and advanced version of Cliff's zero intelligence plus model is introduced in the next section.

3 An Integrated Market and Production Model

The model presented in this paper consists of two parts, the production system and the market system. The production system creates supply and demand that drive the market; the market system generates the prices that drive the production system. First the production system is explained, then the market mechanism. Finally a systematic overview of the complete algorithm is presented.

3.1 The Production System

Suppose we have m activities and n commodities in an economy. The activities can be regarded as production processes, skills or technologies and are denoted by a vector $t_1, \dots, t_m$, whereas the commodities can be regarded as labour, capital, raw materials, or products, and are denoted by a vector $c_1, \dots, c_n$.

We can then formulate an input matrix I and an output matrix O which together describe the production process. In the case of both I and O , rows represent products and columns represent processes. Table 1 provides an example with $m = 4$ and $n = 5$. The columns of the matrices can be read the following way: each of the pairs of columns in Table 1 indicates a process that transforms a combination of input

Fig. 1 An overview of possible network configurations with 25 agents located on grid of 10×10 . The figure on the left displays the agents and their connections when all agents are connected to all other agents. The figure in the middle shows a network with limited connectivity, while the right figure shows limited local connectivity

commodities into a set of output commodities. For example, columns t_1 represent the process that transforms one unit of c_1 into one unit of c_3 , or $c_1 \xrightarrow{t_1} c_3$, and similarly columns t_2 represent the process that transforms one unit of c_1 and one unit of c_3 into two units of c_4 , or $(c_1 + c_3) \xrightarrow{t_2} 2c_4$. Technology t_3 transforms one unit of c_1 and one unit of c_4 into two units of c_5 , or $(c_1 + c_4) \xrightarrow{t_3} (2c_5)$. Technology t_4 transforms one unit of c_1 and one unit of c_5 into $\frac{7}{5}$ units of c_6 and finally $c_6 \xrightarrow{t_5} 2c_1$.

Instead of number subsequent commodities $c_1, c_2, \dots, c_6$ we have chosen to represent commodities by integers: prime numbers for raw materials, and products of integers for composites. This allows us to use the structure offered by the natural numbers, and products literally can be decomposed into their prime factorization to see what raw materials are involved in making them. For example, c_1, c_2 and c_3 are raw materials and assigned as identifiers the primes 2, 3, and 5, and c_4 acquires the number 10 as it is composed of commodities 2 and 5. This is not so relevant for the study in this paper, but in other work new products and technology appear by combining existing products and the natural number structure provides an easy way to name products. In this model, labour is treated explicitly as a commodity (the number 2). Furthermore, there is a special product, which serves as money, and this is the number 3. For the production system in Table 1 this scheme yields the following commodity “names”:

$$\{c_1, c_2, c_3, c_4, c_5, c_6\} = \{2, 3, 5, 10, 20, 40\}.$$

In addition to the technology matrices, the production model includes a set of agents who implement the production process. Agents are located on a grid, i.e. they have x, y coordinates and a z coordinate because there can be more than one agent per location. This grid is introduced in order to allow experimentation in two ways: limiting the agent’s network connections introduces imperfect knowledge, and the grid could also be used to investigate spatial effects (see Fig. 1). However, the simulations in the present paper use the network that connects every agent to every other agent in order to focus on the price mechanism itself.

The agents keep a record of their resources, their skills, their limit prices, and their profit margin. The format and calculation of limit prices and the profit margin is explained further below. The vector to denote the list of skills an agent possesses is a

Boolean vector (e.g. {1, 0, 0, 1, 0}) to indicate the agent has skill 1 and 4), which means that the agent can execute the transformation of columns 1 and 4 of the von Neumann technology matrices. The list of resources is simply a vector to indicate the number of goods in the agent's possession. As mentioned, one of the products (product 3) is designated to serve as the numeraire (a currency) and is used in the exchange of goods, and not, as can be seen in Table 1, in the production process.

All agents are randomly ordered, and the first agent in the list gets to buy, sell, and produce, then the second agent and so on until all agents have had their turn. The order of agents is randomized at each iteration. All of this together constitutes one iteration. The actions of agents are bounded by the following constraints:

- what an agent can buy and sell depends on its resources and on the resources of the other agents in its network;
- no borrowing is allowed, and thus all agents have to maintain a positive balance at all times (this assumption could be relaxed to model business cycles);
- an agent can only execute the skills that it possesses;
- the total number of executions of skills is bounded by a maximum to indicate the total quantity of work one agent can do per time step.

The actions of the agents are profit driven, that is, the agents decide what to produce, and thus which inputs to buy and what to sell, based on what is most profitable to them. In the case profit is not possible, they skip a turn. Profitability is determined by the prices of inputs and outputs, and thus pricing drives the economy. In order to obtain sustained economic activity, it is important that the signals provided by the price levels are appropriately strong and timely. When the signals are too weak or they are generated too late, the cyclical nature of production—from raw materials to intermediate products, to final products, to labour, to raw materials—is disrupted and economic activity grinds to a halt. The algorithm used to set prices and to prevent an economic collapse is explained next.

3.2 Negotiation Rounds

While the production system is in operation, prices must be constantly adjusted to reflect the economic situation. This is not merely a simple convenience; prices need to change constantly to provide appropriate signals to the agents of the production system, since the model, like the real economy, runs without central control. Shortages need to be addressed, and although products have an indefinite shelf life, since resources are scarce, surpluses need to be addressed as well. Prices are the essential signaling mechanism.

Every agent plans what actions are best to take given the production techniques currently in its possession—i.e. it plans to produce those products that are in short supply, and therefore likely to be most profitable, at that moment. At the same time, all agents offer their inventory for sale and plan to stock up on their inputs that threaten to become scarce. This creates a list of plans, one for each agent, of what to buy and what to sell. These plans together imply a supply and a demand curve for the product. The (local) demand and supply curves can be calculated on the basis of this set of

Fig. 2 Supply and demand for product 10. The vertical axis displays the price p , the horizontal axis displays the quantity q . The gray line indicates the downward sloping demand; as prices go down, the demand increases. The black line is the supply curve

plans of all agents in the (local) network, including their limit prices. These curves are not used in the model in the determination of the price, and in fact are not known to the agents, but they are of interest to the modeller and they demonstrate that the prices resulting from the agent's rounds of negotiations more or less converge to the theoretical equilibrium price.

For example, Fig. 2 shows the stepwise demand and supply curve of product 10. This information is not known to any agent. Based on the list of plans, demand and supply quantities can be calculated, as well as the corresponding limit prices. With this information, the supply and demand curves can easily be generated, similar to those in [Smith \(1962\)](#), [Gode and Sunder \(1993\)](#), and [Cliff \(1997\)](#).

Before continuing the description of the negotiations, the format and calculation of the limit prices of buyers and sellers has to be addressed. Every agent has two sets of limit prices, one set of bottom prices, and one set of maximum prices. The former is the lowest price for which the agent is willing to sell. Since no agent wants to sell a commodity for less than the cost to produce it, this is simply the cost price of a product, that is, the value of the inputs per unit of output. The values for all products an agent can produce are updated every time the agent gets its turn in the production iteration described in Sect. 3.1.

The appropriate limit price for buyers is more difficult to determine. How much a producer is willing to pay for one of its inputs depends on the cost of other inputs, as well as the value of the output. It is clear that the total cost of the inputs should be below the value of the output in the case the output is in short supply in order to be able to address that shortage. In the case of a surplus, the situation should be the other way around; the value of the product should be less than or equal to the associated cost to produce it.

First the agent determines a personal price set that reflects the current markets. Products that are in short supply are anticipated to have a price such that their production is profitable; products of which there is a surplus should have a price that makes their production unprofitable. Next, based on this personal price set, the agent determines what it is willing to pay for the inputs. That is to say, products with an anticipated higher price warrant a higher price for the inputs needed to produce the product. The agent first establishes the maximum prices it is willing to pay for the inputs required

for each of the production technologies it possesses. Then the maximum limit price it is willing to pay for an input is the maximum of the maximum prices corresponding to the various production processes in its possession.

There is no unique way to generate such maximum prices. The procedure used in this model first generates a set of anticipated prices that is:

1. close to the current price set;
2. at a level such that cost of input is close to the value of the output offset by 1 such that:
 - (a) in the case of a shortage, the cost of input c is lower than the revenue from the output r :

$$r - c = 1$$

- (b) in the case of a surplus, the cost of input c is higher than the revenue from the output r :

$$r - c = -1.$$

This results in a personal price set that reflects the current markets. Subsequently, the agent uses this price set to calculate how much it is willing to spend on inputs. These prices could be higher than the anticipated prices as long as production remains profitable. Again, this is achieved by determining a price set such that each price is:

1. close to the corresponding personal price;
2. the cost of input is equal to the value of the output where the value of the output is calculated with the anticipated personal price set.

This procedure results in a price set such that if the agent is right in its assessment of the market, it can produce those products which are in short supply with a profit as long as it has to pay less than the limit price for the inputs. Just as for the bottom price for sellers, every time an agent's turn comes around in the production iteration, the agent will update the maximum prices it is willing to pay for inputs.

The agents' plans together with their limit prices and profit margins create a market. An active buyer is an agent with a plan to buy a product. An active seller is an agent that plans to sell its products. However, both buyer and seller attempt to generate as much profit as possible. Buyers will bid less than their maximum price, sellers will ask more than their bottom price, similar to the ZIP traders in [Cliff \(1997\)](#):

$$bid(i, j) = limitprice_{ij} \cdot (1 - \mu_{ij}) \quad (1)$$

$$offer(i, j) = limitprice_{ij} \cdot (1 + v_{ij}) \quad (2)$$

where $\mu_{ij} \in [0, 1]$ is the profit rate of agent j for buying product i and $v_{ij} \in [0, 1]$ is the profit rate of agent j for selling product i .

If prices do not meet, the agents have to adjust their profit margin. The algorithm follows Cliff's work, which is outlined in the previous section and lets agents adapt their profit margin based on the shouts that occur in the market.

Fig. 3 Price negotiations for product 10. The vertical axis displays the price p , the horizontal axis displays the negotiation steps t . Light gray squares indicate bids, black dots indicate offers, and gray diamonds are the prices that buyers and seller have agreed upon. In order to make a possible trend visible, the gray diamonds have been connected. The horizontal line between the black triangles shows the equilibrium price (Eq.)

Fig. 4 Price negotiations for product 2 with an almost fully elastic supply curve. In both graphs the vertical axis displays the price p , in the left graph the horizontal axis displays the quantity q and in the right graph the horizontal axis displays the subsequent steps in the negotiation process. In the graph on the left, the gray line indicates the downward sloping demand. The black line is the supply curve. In the graph on the right light gray squares indicate bids, black dots indicate offers, and gray diamonds are the prices that buyers and seller have agreed upon. The equilibrium price (Eq.) and quantity where supply and demand meet ($p \approx 0.86$ and $q = 14$) can be seen in the right graph by the deal prices which are converging to 0.86 from above, and the number of completed deals, i.e. the number of gray diamonds (14). Different bid levels are caused by different maximum prices that go with different skill sets

The negotiation dynamics that result from this price mechanism are illustrated in Figs. 3, 4, 5 and 6.

Figure 3 shows the negotiation results corresponding to the demand and supply curves of Fig. 2, that is, the limit prices of the negotiations in Fig. 3 are the ones that generated the demand and supply curve of Fig. 2. As can be seen in Fig. 2 the demand and supply are balanced, and the curves intersect at $q = 14$ and $p \approx 1$. Figure 3 shows the shouts of the corresponding negotiation rounds: light gray squares indicate bids, black dots indicate offers, and gray diamonds indicate when an offer and a bid cross. In the case the last shout was a bid, the lowest offer of all offers below the bid is accepted and a transaction takes place. In the case the last shout is an offer, the highest

Fig. 5 Price negotiations for product 10. An example of a seller's market. In both graphs the vertical axis displays the price p , in the left graph the horizontal axis displays the quantity q and in the right graph the horizontal axis displays the subsequent steps in the negotiation process. In the graph on the left, the gray line indicates the downward sloping demand. The black line is the supply curve. In the graph on the right light gray squares indicate bids, black dots indicate offers, and gray diamonds are the prices that buyers and seller have agreed upon. The equilibrium price (Eq.) and quantity where supply and demand meet ($p \approx 1.28$ and $q = 8$) can be seen in the right graph by the deal prices which are converging to 1.28 from below, and the number of completed deals, i.e. the number of gray diamonds (8)

Fig. 6 Price negotiations for product 20. In both graphs the vertical axis displays the price p , in the left graph the horizontal axis displays the quantity q and in the right graph the horizontal axis displays the subsequent steps in the negotiation process. In the graph on the left, the gray line indicates the downward sloping demand. The black line is the supply curve. In the graph on the right light gray squares indicate bids, black dots indicate offers, and gray diamonds are the prices that buyers and seller have agreed upon. The equilibrium price (Eq.) and quantity where supply and demand meet ($p \approx 1.05$ and $q = 15$) can be seen in the right graph by the deal prices which are converging to 1.05, and the number of completed deals, i.e. the number of gray diamonds (15)

bid of all bids above the offer will be the transaction price. There are 14 transactions (as expected) and these center around $p = 1$ as expected.

Figure 4 gives an example of an elastic supply curve; each seller has practically the same lower bound. Unlike the experiment of chart 4 in Smith (1962, p. 119), where in the classroom experiment done by Smith the price converged to a level slightly higher than the equilibrium, here the transaction prices do converge to the theoretical equilibrium. The difference can be explained by the fact that the negotiation simulation of Fig. 4 is a case of a buyers' market instead of the balanced market in Smith (1962). Initially the price is driven up, but since there is much more supply than demand, ultimately the price is reduced to equilibrium levels.

Figure 5 shows an example of the opposite situation: a sellers' market where there is much more demand than supply. Sellers have floor prices roughly between 0.9 and

Fig. 7 The creation of new prices

1.1, but the sales that occur are mostly at substantially higher prices and keep going up until all eight units of product 130 offered are sold.

Figure 6 shows a similar unbalanced situation—far more demand than supply—except that now most buyers are in the market with a maximum price below the minimum price of sellers, a fact of which the market participants are unaware. The limited number of buyers with a higher limit price find it easy to reach a deal somewhere between their maximum price of 1.5 and the sellers minimum price of approximately 1. The other buyers are forced to increase their bids, but to no avail, as even at their limit prices they still do not find a supplier.

3.3 The Price Mechanism and Production System Algorithm

This section combines the price mechanism of Sect. 3.2 and the production model of Sect. 3.1 into an integrated representation of the dynamics of a market economy.

As Fig. 7 shows, the generation of new prices begins with all agents making a private plan as to what is best for them to do given their knowledge of the current state of affairs. The agents take into account the current stock levels to see what should and what should not be produced. They also consider the resources that are currently available to them, either because they can be acquired on the market or because they are already in their possession. The plans are not executed, but are taken to market. In addition to the number of units that agents want to buy or sell, the agents also bring their limit prices—minimum prices for sellers and maximum prices for buyers—and their profit margins. Then the negotiations begin, and for each product for which there is a supply and demand a double auction is held, resulting in transactions and new prices.

The scheduling of the actual negotiations is illustrated in Fig. 8. For each of the products in Fig. 7 for which there is a market, the loop of Fig. 8 is executed. The active buyers and sellers for a commodity are randomly ordered and the first agent is added to the market. This means that a buyer enters with a bid or a seller enters with an offer as given by Eqs. 1 and 2. Then the next agent is added to the market and so on. When an agent revises its shout (bid or offer), the old shout of this agent is removed from the market, and it is checked if bids and offers cross. If so, the transaction takes place and the buyer and seller are removed from the market. With every shout all agents adapt their profit margin as indicated in the pseudo code in Sect. 2. This procedure of shouting bids and offers is executed until all agents on the list have entered the mar-

Fig. 8 The agents' negotiation module. Agents enter the market in a random order. When all agents have announced their bid or offer, and there are still agents that haven't made all their desired purchases and sales, agents revise their former bid or offer. The agents do so in a new random order to achieve an unbiased entry procedure so that none of the agents benefits from a particular entry order

Fig. 9 The scheduling of all steps that together constitute an iteration

ket. Some agents will have reached a deal, others will have had to adjust their profit margin and will have not reached a deal yet. For those who still want to trade a new randomly ordered list of active agents is created and these agents revise their entries on the market while continuously adjusting their profit margin with every new shout. This continues until there are no more active buyers or sellers, or until the remaining agents have reached their limit prices and no further deals are possible. The new price is the price of the last successful exchange.

The last flow chart (Fig. 9) puts the previous two figures into context. It shows the scheduling of steps that have to be taken in order to execute one iteration of the model. All agents are ordered randomly in a list, and every agent gets its turn to buy, sell and produce with the goal to maximize its profits. Before an agent can do so, actual prices need to be set, and this is achieved by organizing a CDA for all products for which there is both demand and supply. This is an imaginary auction; the negotiations and exchanges only take place to have prices of consecutive exchanges converge to a new price that reflects the current state of the market. Once new prices are set, the agent gets to buy, sell and produce products, and to update its limit prices—both floor prices for sales and maximum prices for purchases. Then the next agent in the list gets its turn, but before it can buy, sell and produce new prices are set by organizing

imaginary CDA's and so forth and so on. Once all agents in the list have had their turn, the iteration is complete. A simulation consists of many iterations. The results of a simulation run—time series of prices, stocks, production levels, and profitability—are discussed in the next section.

4 Results

So far one can observe, the market (or price mechanism) is capable of generating prices which correspond to the underlying supply and demand curves, but the bigger task of the market is to organize both allocation of resources and evolution of the production system (Potts 2001). This is why the evolving production system in Straatman et al. (2008) was coupled to an agent-based market. This paper focuses on a new market mechanism that will be further used to study evolving economies too, but for this paper experimentation is limited to a fixed production system. Nevertheless, the market still has to function as a mechanism to organize allocation of scarce resources. In order to do so, the market has to provide the production system with timely and clear signals, so that the production of all required commodities occurs on a continuing basis. To verify if the new pricing mechanism based on the actual negotiations by agents (instead of the Walrasian auctioneer of the earlier version of the model in Straatman et al. (2008)) is capable of orchestrating an appropriate allocation of resources, the market is coupled to the production system and the resulting production activity is analyzed. Model validation consists of a comparison between generated activity and predicted activity, where the predicted production activity is based upon the Leontief inverse of the input-output matrix (that can easily be constructed from the pair of von Neumann technology matrices) and is used to calculate activity to generate external demand.

4.1 Initial Settings

The simulations in the present paper use the von Neumann production system with products 2, 3, 5, 10, 20, 40 as outlined in Table 1 and a random distribution of 25 agents that are all connected to one another as displayed in the left figure of Fig. 1. The number of agents has been chosen such that it is not too large, which would lead to long simulation times, and not too small. A too small number would make it difficult to distribute the different skills required in the production cycle. Each agent obtains a random boolean vector of length five representing its skills, for example agent $\{1, 8, 1\}$, $\{0, 0, 1, 1, 0\}$ implies that agent $\{1, 8, 1\}$ can execute the transformations of the von Neumann technology matrices in columns three and four. It is made sure that every technology is present in the model as well as that each agent has at least one skill. Corresponding to the skills it possesses, each agent obtains the same abundant amount of the numeraire and a limited quantity of resources that allow the agent to produce in case that provides financial gains. Whether or not production is profitable depends on the price of input and output, and at the start of a simulation the prices are set equal to a random perturbation of a feasible price set. Feasible here implies that there exist prices for each of the commodities such that the produc-

tion system can be operated with a profit, in other words there exist prices $p_1, \dots, p_6$ such that

$$(O - I) \cdot \{p_1, p_2, p_3, p_4, p_5, p_6\} \geq 0,$$

where O is the output matrix and I is the input matrix of Table 1. von Neumann (1946) provides the conditions under which a feasible price set exists. It can easily be verified that $\{1, 1, 1, 1, 1, \frac{10}{7}\}$ is such a price set. Similarly, the agents' limit prices for buying and selling are initially equal to this random perturbation of this feasible price set. The profit margin for each agent is set equal to $\frac{1}{5}$, meaning that at each start of an imaginary auction, organized to determine current prices of goods, each agents attempts to buy (sell) at a price 20% below (above) its limit price. This configuration allows the agents to start production, to repeatedly adapt limit prices based on the actual economic situation, and to repeatedly set new prices by negotiation in order to both maximize personal gain and to allocate resources effectively.

Obviously the null-state is a stable state for a simple production system: no production leads to no demand and therefore no shortage. In order to rule this state out, the agents in the model are confronted with depreciation or "wear and tear", and as such they are forced to use some of their resources for consumption or for maintenance, or for unavoidable fixed costs. This concept is implemented by an exogenously fixed demand for one of the products (*product2*) by each agent at each time step (comparable to a metabolic rate in biology), independent of production.

The artificial economy here is closed, so there is no external demand as such, but the depreciation and forced expenditures on maintenance generates such external demand. The cost of depreciation should not be so high as to drain resources out of the otherwise closed system faster than can be supported by the collective production system. On the other hand, it should also not be so low as to allow the production system to quickly accumulate large surpluses that allow the system to be passive most of the time and only active some of the time, although the biggest problem with this scenario is the increase in simulation run time with very little action. Based on experimentation with different metabolic rates the simulations here are run with a metabolic rate of $\frac{1}{2}$ units of the first product to keep the agents alert. The production output levels per iteration required to generate sufficient "food" are:

$$\{12.5, 0, 0, 0, 0\} \cdot (O - I)^{-1} \cdot O \approx \{43.8, 3.9, 7.8, 15.6, 21.9\}$$

where $\{12.5, 0, 0, 0, 0\}$ is the vector that indicates the exogenous metabolism of all agents (25 in this case) at a metabolic rate of $\frac{1}{2}$ units of the first product, O is the output matrix and I is the input matrix. $\{12.5, 0, 0, 0, 0\} \cdot (O - I)^{-1}$ gives the activity levels required to generate the exogenous demand, and the product of these activity levels and the output matrix gives the expected output per iteration in order to sustain the population given the specified metabolic rate.

4.2 Results Analysis

Since no agent is capable of producing every product, the raw output data per agent shows wildly fluctuating production numbers varying from 0 to the maximum output. But also the output per iteration fluctuates, which can be partly explained by the small number of agents in the system and by the random ordering of agents in the execution of an iteration—i.e. an agent can get its turn before its suppliers have had their turn. However, when the data are displayed as a moving average over a window of 250 agents' actions, which corresponds to 10 iteration since this simulation had a total of 25 agents, a much better idea of the dynamics of the output is obtained. Figure 10 shows the output as a moving average over ten iterations.

One can observe that the levels are slightly above the expected levels based on the Leontief inverse. This is an indication that the price signals are more or less appropriate, but not perfect. In the case of a production system where resources are scarcer than the case under review here, and thus where mis-allocation of resources will lead to a halt of the production process due to shortages from which the collective system cannot recover, the price mechanism as it is used here would not suffice. Whether or not this can be remedied by different parameter settings or requires more advanced trading strategies is unclear at the moment.

In order to get insight into the coupling between price on the one hand and supply and demand on the other, it is necessary to look at the effects of changing prices and supply and demand. Since the agents here are driven by stock levels and profitability the question is what happens to stock levels and profitability. Figure 11 displays the moving average of production numbers, stock levels, prices and profitability.

Figure 11 shows the interaction between price and volume of production and *visa versa*. The top row shows those production averages and the second row gives the resulting stock levels for each of the products. This information is available to the agents and based on these stock levels, agents make their plans and start the negotiations as explained in Sect. 3. The prices that follow from these negotiations are displayed in the third row. Finally, the bottom row shows the moving average of profitability. A production process for a product is either profitable (+1) or unprofitable (−1). It is profitable when the input costs are lower than the value of the output, and unprofitable when the cost of the input is higher than the value of the output. By taking the average over a window, we obtain values in the $[-1, 1]$ interval, where a value of 1 indicates that the production process has been profitable for at least 250 consecutive time steps, and 0 means that in the window an equal number of time steps were profitable and unprofitable. It is the profitability information that agents use to decide on what to produce and what not, and thus the top row with production data is a direct result of the bottom row with profitability data. Conversely, as just explained, the bottom row is a result of production and stock levels. When the stock levels are too low, the prices should go up and the production should become profitable. The graphs show something similar, but not so explicitly. Just as in the real economy, an increase in the price of a commodity does nothing for the profitability of the production process of the commodity when the cost of inputs increases at a higher rate.

Regardless of the less than clear profitability signals resulting from the prices, the artificial economy does stay afloat. Based upon a large number of experiments

Fig. 10 Average production output per 10 iterations. Although the average production levels fluctuate, they do this around the expected levels as calculated by the Leontief inverse

Fig. 11 The feedback loop of production (*top row*), stock levels (*second row*), pricing (*third row*) and profitability (*bottom row*) illustrated. In each of these graphs, the raw data are transformed into moving average data over 10 iterations in order to improve the readability

(more than 100) with numerous variations on price mechanisms by the authors, it can be said that this is no small feat. The fact that the negotiating agents agree on prices close to the theoretical equilibrium and that these prices suffice to keep such a small system—hence one without much room for error—functioning suggests that the price mechanism is effective. Together with the model of the production system it provides again an alternative look on the invisible hand, as a highly dynamic system is run without central command. Nevertheless there is room for improvements.

5 Conclusions

The price mechanism used here is the simple version of [Cliff \(1997\)](#)'s zero intelligence plus traders where the agent's trading behaviour is controlled by eight parameters. In the meantime there is an advanced version that uses sixty parameters to control the trading strategies ([Cliff 2006](#)). This large number of parameters requires a genetic algorithm to optimize the parameter settings. Further experimentation will have to show if the advantages of the improved trading strategies warrant the substantially larger parameter set of the expanded model.

What makes the present model stand out is the ability to provide a complete picture of both production and market system in one model. The immediate advantage is that the model can cope with multiple products that are interrelated, instead of the single commodity markets that are common in most approaches. More generally, the model can be considered to be an agent-based implementation of a general equilibrium model. As such, it can be used to gain additional insights into the nature of conventional microeconomic models. In particular, this model is characterized by heterogeneous agents with shifting production plans. Furthermore, the agents have bounded rationality due to a simple set of rules to determine their limit prices, profit margins and actions. The fact that such a model can nevertheless attain an equilibrium set of prices and outputs—one that is essentially in agreement with the vector of outputs given by the Leontief inverse—suggests that conventional microeconomic models may be relatively robust with respect to various market imperfections. Further experiments, focused on the spatially defined trading networks, can shed additional light on the impact of market imperfections as well as on the limiting conditions for the appearance of a single general equilibrium. For example, it may be expected that if trading networks are too restricted, regional markets may appear, each with its own set of prices.

In addition the model allows additional products and technology. It will be very interesting to see how prices develop as novel elements are included, but that is work for later. What is certain is that the new price mechanism is much more flexible than the earlier version, and thus that the improved model of an evolving artificial economy (and real ones too) might generate more interesting results.

Acknowledgments Funding for this research was provided by the Institute for Sustainable Energy, Environment and Economy (ISEEE) at the University of Calgary.

References

- Canterbery, E. R. (2001). *A brief history of economics: Artful approaches to the dismal science*. Singapore: Word Scientific Publishing.
- Chaggar, S. S., Noble, J., & Cliff, D. (2008). The effects of periodic and continuous market environments on the performance of trading agents. In S. Bullock, J. Noble, R. A. Watson, & M. A. Bedau (Eds.), *Artificial life XI: Proceedings of the eleventh international conference on the simulation and synthesis of living systems* (pp. 110–117). Cambridge: MIT Press, MA.
- Chen, S.-H. (2000). Toward an agent-based computational modeling of bargaining strategies in double auction markets with genetic programming. In K. Leung, L.-W. Chan, & H. Meng (Eds.), *Intelligent Data Engineering and Automated Learning-IDEAL 2000: Data mining, financial engineering, and intelligent agents* (pp. 517–531).
- Chiarella, C., & Iori, G. (2002). A simulation analysis of the microstructure of double auction markets. *Quantitative Finance*, 2, 346–353.
- Cliff, D. (1997). Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical report, School of Cognitive and Computing Sciences, University of Sussex.
- Cliff, D. (2006). Zip60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms. *IEEE Transactions on Evolutionary Computation*.
- Cliff, D., & Preist, C. (2001). Days without end: On the stability of experimental single-period continuous double auction markets. Technical report, Information Infrastructure laboratory, HP Laboratories Bristol.
- Das, S. (2003). *Intelligent market-making in artificial financial markets*. Master's thesis, Massachusetts Institute of Technology.
- Davis, D., & Holt, C. (1993). *Experimental Economics*. Princeton, NJ: Princeton University Press.
- Friedman, D. & Rust, J. (Eds.). (1993). *The double auction market: Institutions, theories, and evidence* (Vol. XIV). New York: Addison-Wesley.
- Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1), 119–137.
- Kauffman, S. A. (2000). *Investigations*. Oxford: Oxford University Press.
- LeBaron, B. (2001). A builders guide to agent based financial markets. *Quantitative Finance*, 1(2), 254–261.
- LeBaron, B., Arthur, W., & Palmer, R. (1999). Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23, 1487–1516.
- Ma, H., & Leung, H. (2008). *Bidding strategies in agent-based continuous double auctions*. Basel: Birkhauser.
- Potts, J. (2001). Knowledge and markets. *Journal of Evolutionary Economics*, 11(4), 413–431.
- Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2), 111–137.
- Straatman, B., White, R., & Banzhaf, W. (2008). An artificial-chemistry based model of economies. In S. Bullock, J. Noble, R. A. Watson, & M. A. Bedau (Eds.), *Artificial life XI: Proceedings of the eleventh international conference on the simulation and synthesis of living systems* (pp. 592–599). Cambridge: MIT Press.
- von Neumann, J. (1946). A model of general economic equilibrium. *Review of Economic Studies*, 13(1), 1–9.