Data-Driven Agent-Based Simulation of Commercial Barter Trade

Peter Haddawy, Khaimook Dhananaiyapergse,
Yongyos Kaewpitakkun, Thai Bui
CSIM Program, Asian Institute of Technology
haddawy@ait.ac.th, khaimook@chula.com, yongyosa@tot.co.th, bqthai@cit.ctu.edu.vn

Abstract

In this paper we present TRADES, a data-driven agent-based simulator for barter trade exchanges. We provide an overview of the barter trade exchange industry, focusing on the operational aspects of trade exchanges and motivating the design of our simulator. Our simulator is built by learning probabilistic models of company purchase behavior using transaction history data from an operating trade exchange. We quantitatively evaluate the accuracy of our simulator by comparing simulated trade to the transaction data, showing a high degree of agreement between the two. We also demonstrate use of the simulator to evaluate the effectiveness of a particular trade brokering strategy.

1. Introduction

Agent-based simulation has become a popular method of evaluating the effectiveness of trading agent and market mechanism design and for observing the effects of various trading behaviors on market parameters. It has proven to be particularly useful in situations where theoretical analysis either makes too many simplifying assumptions or is simply not possible due to problem complexity [1]. It has been used, for example, to evaluate the effectiveness of auctions [1], [2], [3], buyer coalition schemes [4], and brokered trade [5]. In previous work on agent-mediated marketplace simulation, agent behavior is either explicitly specified in terms of rules or parameters such as value functions, or it is initially specified and then allowed to adapt through various learning mechanisms. This approach makes verifying the accuracy of the simulation difficult, so that when accuracy has been addressed, it has been addressed qualitatively by specifying particular types of participating agents and/or governing rules and then comparing the market dynamics to economic theory or to informal observations.

In contrast, in this paper we take a data-driven approach to agent-based marketplace simulation. We learn agent trading behavior from marketplace transaction history data. This enables us to then *quantitatively* evaluate the accuracy of the simulation by comparing the simulated trading behavior to the actual transaction history using standard evaluation techniques from simulation and machine learning.

The model used in this paper is that of the barter trade exchange, also called retail or commercial barter. A barter trade exchange is a collection of businesses that trade their goods and services, managed by an intermediary. We call the collection of businesses the *barter pool* and call the intermediary the *trade exchange*. In modern barter trade exchanges, businesses do not exchange goods directly in the bilateral fashion of traditional barter. Rather, modern barter is multilateral, using a form of private label currency. The trade exchange issues trade dollars to the member businesses and acts as a neutral third party record keeper. When a company sells a good, they receive credit in trade dollars, which they can then use to purchase goods from other members. The value of the trade dollar is tied to the US dollar by not permitting businesses to charge more for their goods in terms of trade dollars than they do in US dollars in the open market, thus preventing devaluation of the currency.

Commercial barter is an attractive area for simulation and for experimenting with market design mechanisms because a barter pool is a relatively closed economy about which we have very detailed information due to the book-keeping function of the trade exchange. The trade exchange maintains a general profile for every member business, as well as complete records of all transactions between members. A barter pool has many similarities with a traditional economy, with the trade exchange playing a role analogous to that of the federal government in regulating the economy. The exchange controls such variables as monetary supply, interest rate, rate of commission (analogous to revenue tax), and even supply and demand through its ability to selectively recruit new member businesses. Interestingly, although it has control over all these parameters, the trade exchange works to stimulate the barter pool economy primarily by making referrals to member businesses through trade brokers.

The success and survivability of the barter business add to its attractiveness as a model to study. The barter trade exchange industry has existed for over forty years, surviving numerous changes in the economic landscape. The International Reciprocal Trade Association [6] estimated that the total value of products and services bartered by businesses through barter companies reached USD 7.87 billion in 2001. There were an estimated 719 trade companies active in North America in 1999 with some 471,000 client businesses [7]. Examples of active barter trade exchanges with a Web presence include BizXchange, ITEX, BarterCard, and Continental Trade Exchange.

The rest of this paper is organized as follows. In Section 2 we provide a description of the operation of barter trade exchanges. In Section 3 we provide an overview of our Barter Trade Exchange Simulator (TRADES), describing its functionality, parameters, and overall architecture. In Section 4 we describe the implementation of the TRADES system, covering the learning of company purchase behavior models and the algorithms used to generate supply and demand. In Section 5 we describe the results of our empirical evaluation of the simulator. Section 6 covers related work and Section 7 presents conclusions and directions for future research.

2. Commercial Barter

Given its important role in B2B commerce, there is a surprising lack of literature on the barter trade exchange industry. An exception is the work of Cresti, which examines theoretical economic rationale for development of the barter industry in industrialized countries [8], as well as investigating the macroeconomic variables influencing the industry in the United States [9]. But there exists no formal literature describing the barter trade exchange industry on an operational level. Our interest lies in understanding the parameters governing trade exchanges as well as how managers and brokers manage the operations of the exchange in order to maximize their company's profits. Thus our first step in conducting this work was to gather information through extensive interviews with industry experts. We also communicated with them periodically to verify the assumptions behind our models. We interviewed two executives at BizXchange (www.bizx.bz), a relatively new but rapidly growing trade exchange located in the San Francisco Bay and Seattle areas. Since its inception in January 2002, BizXchange has grown to include over 600 member businesses. The two executives we interviewed have over 28 years of combined industry experience, have founded and built several successful barter networks, and have served on the Boards of the International Reciprocal Trade Association and the National Association of Trade Exchanges.

A barter pool can be viewed as a carefully managed small-scale economy. Managers of trade exchanges attempt to recruit member businesses in such a way that supply and demand for each product category in the pool are approximately balanced[1]. Member businesses are typically small to medium size enterprises that offer products and/or services. They fall into the broad categories of operating expenses, employee benefits, and travel and entertainment. Henceforth we will use the term *goods* or *products* to refer to goods and services.

It is a common misconception that the primary benefit of barter is to avoid taxes. In fact, the US Tax Equity and Fiscal Responsibility Act, passed in 1982, legislated that barter income be treated as equivalent to cash income and taxed on the same basis. Cresti [9] shows empirical evidence that barter is adopted to increase profits and gain a competitive edge and that barter is, in fact, complementary to the cash economy.

When a business joins a trade exchange, it typically pays a membership fee. This represents a small fraction of the revenues of the trade exchange. The primary revenue is made by charging a fee to the buyer and seller on each transaction. The fee is typically in the range of 6 - 7.5% and is payable in US dollars. When a business joins the trade exchange, they are issued a line of credit in trade dollars, which permits them to make purchases without first having to sell and also gives them flexibility in conducting transactions. The trade exchange charges interest on negative balances, usually at the same rate as major credit cards. In order to give a company some control over how much of their profits are accrued in terms of trade dollars, the trade exchange permits the member to set an upper limit on the amount of trade dollars they are willing to accumulate. The credit line and upper limit define the *financial operating range* of the business within the barter pool.

Each member is assigned to a trade broker. A broker typically represents a set of 150 - 200 client businesses. The broker's job from the standpoint of the client is to help the client sell his goods and to inform him of goods he might like to buy. The broker's job from the standpoint of the trade exchange is to stimulate trade, since the exchange's revenues are directly tied to trade volume. The broker stimulates trade by working to help clients spend their trade dollars when they have positive balance and generate sales when they have negative balance. The broker's primary tool is the referral, referring potential buyers to suppliers. Note that member businesses are under no obligation to follow the broker's referrals, but experience from trade exchanges shows that they

[1] Although managers *attempt* to keep supply and demand balanced, it is not the case that they are, in fact, balanced at any point in time. Therefore we do not assume that supply and demand are equal or even near-equal.

generally do. While the goods a business has to sell are stated explicitly, those that the business wants to buy may be explicitly stated or may be predicted by the broker based on things like the type of business and other goods that the business has purchased in the past.

3. TRADES System Overview

The TRADES system is designed to simulate barter trade so that the effectiveness of various brokering strategies can be evaluated under different barter pool compositions and different member business parameters. We start by describing the assumptions made in designing the simulator. We assume that trade occurs in business cycles: first businesses' supplies and demands are determined, the businesses are supplied with referrals, the businesses act on the referrals, and the cycle repeats. Supply and demand in each cycle can be represented by a *requirements matrix* in which each row represents a member business, each column represents a category of goods, and matrix entries represent quantities to buy or sell. Each business can buy multiple categories of goods and we make no assumptions about the relationship between supply and demand in the barter pool. The internal state of each business is characterized by its current balance, its credit line, and an upper limit on allowed balance. Once a business reaches its credit limit, then it can no longer buy without selling, so its row in the matrix will show no demand. Similarly, when a company reaches its upper limit, it can no longer sell without buying, so its row will show no supply. Figure 1 shows an example requirements matrix for a barter pool with four companies and four product categories. Company C1 has $500 of product P1 to sell and wishes to buy $100 each of products P3 and P4.

	Product			
Company	**P1**	**P2**	**P3**	**P4**
C1	+500	0	-100	-100
C2	-200	-200	+400	0
C3	0	+100	0	-200
C4	-200	+100	0	-300

Figure 1. Example requirements matrix

The TRADES system architecture is shown in Figure 2. First companies to include in the simulation are created. Then the basic simulation cycle begins by generating supply and demand, represented as the requirements matrix. Since TRADES is designed to experiment with various brokering strategies, it includes an interface to an external brokering module. TRADES provides the requirements matrix as input to the brokering module and receives back a set of referrals, matching buyers with sellers. Each company then acts on the referrals based on the company behavior parameters discussed below. This results in new company balances, and the cycle repeats.

TRADES contains two basic groups of functions: company creation, and company parameter specification. The system stores a set of probabilistic models that specify company purchase behavior, categorized by the product category the company sells. How the models are learned from data is described in Section 4. The user can choose to create instances of companies to include in the barter pool. The user can either specifically select individual company models or indicate a number of companies to create from a category. In the latter case, TRADES randomly selects the models to use. It is possible to create more than one instance of a given purchase behavior model. This does not mean that the two created companies will necessarily have the same behavior since company behavior is determined by other parameters as well. Figure 3 shows an example of creating companies from the learned models.

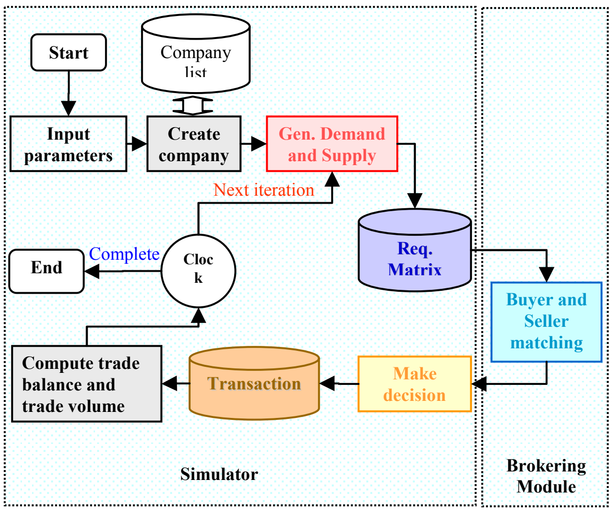

Figure 2. TRADES system architecture

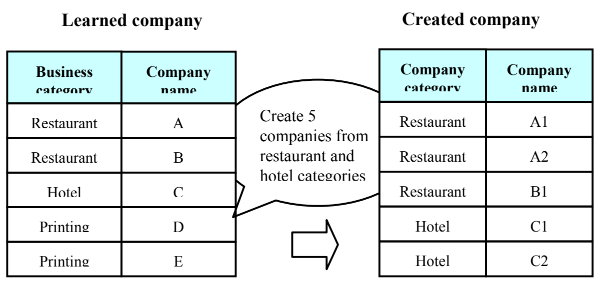

Figure 3. Example of creating companies

For each company created, the user can specify the following attributes: starting trade balance, financial operating range (credit limit and maximum balance), probability of following a referral, probability to buy out of the barter pool, and probability to buy if a referral matches an unfilled product need from the previous week.

A simulation run is launched by simply specifying the number of weeks for the simulation. The simulator can display data on individual company and overall barter pool behavior over time, including individual company balance, individual company trade volume, barter pool absolute balance, and barter pool trade volume. The barter pool *absolute balance* is defined as the sum of the absolute values of the balances of the companies in the pool. It is a measure of the balance of trade in the barter pool [14]. An example of the TRADES output screen is shown in Figure 4.

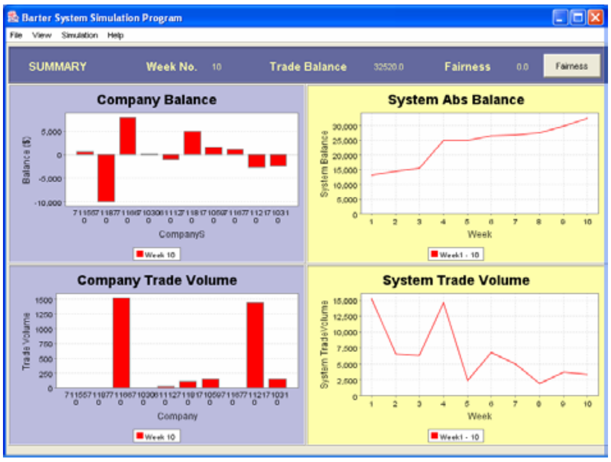

Figure 4. TRADES output screen

4. Implementation

4.1 Learning Purchase Prediction Models

We obtained 16 months (66 weeks) of transaction history data from BizXchange. The data specified the buyer, seller, date, amount, and product category for each transaction. Each member business supplied only one product category. The data included each company's credit line, which ranged from \$500 to \$20,000, and the company's upper bound on balance, which ranged from \$10,000 to \$50,000. The number of companies started at only 4 in week one and steadily grew to 264 by week 66. The total number of transactions was 1,887. The average number of suppliers per product category in which there was at least one buyer was 4.8 and the average number of buyers per category was 1.9.

To prepare the data for learning of company purchase prediction models, we filtered out companies that had less than 16 purchases, i.e. on average one per month. After filtering, we were left with 1,535 transactions from 46 companies in 26 product categories. We experimented with a number of different Bayesian network models to see which one would yield the best predictive accuracy. We tried learning five structurally different naïve Bayes clustering models [10] with various numbers of states of the clustering variables. We also tried a simple naïve Bayes model. We evaluated each model by training on the first 70% of the data and then testing on the remaining 30%, using ROC analysis [11] to measure model predictive accuracy. ROC analysis is a standard method for evaluating the predictive accuracy of probabilistic models. The naïve Bayes clustering models had ROC values ranging from a low of 0.71 to a high of 0.78. The simple naïve Bayes model had the highest ROC value of 0.82.

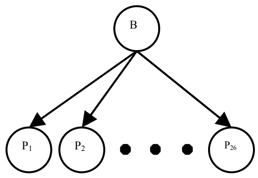

Figure 5. Naïve Bayes purchase prediction model

Thus we chose the simple naïve Bayes model shown in Figure 5 for use in our experiments. The model contains one node (B) with the identifiers of the businesses and one node (P_i) for each of the 26 product categories, indicating the probability that the given business buys that product in any given week. For simulation we also require a prediction of the dollar amount of each product that a business will purchase. It would be possible to have the states of the product nodes correspond to purchase amounts, but this would result in a model of too high a dimensionality relative to the data. Thus to predict product purchase amounts, we created a histogram for each company, dividing the purchase amount into intervals of ten dollars.

4.2 Determining Demand and Supply

We use Monte Carlo simulation, to determine company product demand. First we determine which products each company wishes to buy using the inverse-transform method for the discrete case [12]. For each company and each product category in the barter pool, we generate a uniformly distributed random variable between 0 and 1. If the generated value is less than the purchase probability for that product category, then we say that the company wishes to buy that product in that time period. We then similarly determine the purchase amount for each product by generating another random variable and taking the midpoint of the indicated quantity interval from the histogram.

Since each company supplies only one product category, supply is simply determined by taking the difference between the company's current balance and the upper bound of its financial operating range. This results in a complete requirements matrix for the time period, which can then be passed to the brokering module.

5. Evaluation

We ran two different sets of experiments. The purpose of the first set was to determine the accuracy of the simulation, while the purpose of the second was to evaluate the effectiveness of a particular brokering strategy in stimulating trade.

We evaluated the accuracy of the simulator by comparing simulated barter pool trade with that from the BizXchange transaction history. The simulation was run as follows. The product demands and supplies were determined by Monte Carlo simulation, as described above. Since we wanted to compare trade behavior with that in the barter pool, we did not use a brokering module to match buyers and sellers. Doing so would introduce an exogenous influence on the behavior. Rather, we used the suppliers that each company had gone to in the transaction history data in order to determine the buyer/seller pairing. For example, suppose that the Monte Carlo simulation indicates that company C1 needs to buy $300 worth of carpet cleaning services in week 20. We look in the transaction history to see which suppliers C1 ever purchased carpet cleaning services from and choose one of them with a probability proportional to the amount purchased from that supplier in relation to the other carpet cleaning service suppliers that C1 purchased from.

In the transaction history data, the number of companies grows from 4 in the first week to 263 in the last week. Accordingly, we ran our simulation so that it added the appropriate number of companies in the appropriate product categories each week. We ran the simulator for periods of 100, 200, 500, and 1000 weeks, each ten times. We evaluated the accuracy of the simulation in terms of four parameters: purchase amount and sales amount averaged over all companies, and the absolute trade balance and the overall trade volume of the entire barter pool. The last two parameters were averaged over the number of weeks in the simulation run, while the first two were averaged only over the number of weeks that each company was a member of the barter pool. As our measure of accuracy, we used the efficiency index [13], which measures the agreement between simulated and actual values of a given parameter (e.g. trade volume) as a proportion of the total range of that parameter in the data. The efficiency index (E) is defined as

$$E = 1 - \frac{SE}{ST}$$

$$\text{where} \quad SE = \sum_{i=1}^{N} (y_i - \hat{y}_i)^2 = \text{sum square of errors}$$

$$ST = \sum_{i=1}^{N} (y_i - \bar{y})^2 = \text{total variation}$$

$$y_i = \textit{observed value}$$

$$\hat{y}_i = \textit{simulated value}$$

$$\bar{y} = \frac{1}{N} \sum_{i=1}^{N} y_i$$

$$N = \textit{number of data points used}$$

The value of the efficiency index ranges from a maximum value of 1 to a minimum of $-\infty$. Results of these experiments are shown in Table 1 and Figure 6. Since this is a Monte Carlo simulation, the accuracy increases asymptotically as we increase the length of the simulation. At only 200 weeks, the efficiency index of all parameters is already 0.90 or above, showing excellent agreement with the data.

Table 1. Efficiency index

Aspects	Efficiency Index			
	100 wks	200 wks	500 wks	1000 wks
Purchases	0.71	0.90	0.93	0.93
Sales	0.74	0.90	0.94	0.93
Absolute trade balance	0.86	0.93	0.94	0.92
Trade volume	0.77	0.93	0.96	0.95

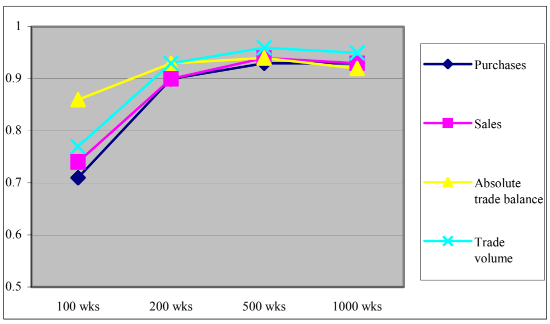

Figure 6. Efficiency index

Our second experiment evaluated the effectiveness of a brokering module that emulates the practice of trade exchange brokers by matching buyers and sellers in such as way that single-period trade volume is maximized, while balance of trade is maintained as much as possible. Details of the brokering algorithm can be found in [14]. We were interested in seeing whether the strategy of maintaining trade balance helps to maximize trade volume over the long run. Each simulation used the same set of 130 companies and 26 product categories, with 3 - 7 suppliers per product category. Company balances all started at zero trade dollars. Company financial operating ranges were set to be identical to those in the BizXchange barter pool. We ran ten simulations of 100 weeks each. In one set of simulations we matched buyers and sellers so that trade volume in that week was maximized and absolute balance was minimized. Absolute balance is defined as the sum of the absolute values of the balances of all companies in the barter pool. In the other simulation, we maximized trade volume in each week, ignoring balance. The probability for each company to follow the brokering module's referrals was set to one. The results of the simulation are shown in Table 2. Use of balance optimization results in a reduction in absolute balance of 24%, with an accompanying increase in trade volume of 40%. These results provide support to the rule of thumb used by trade brokers that maximizing single-period trade volume while maintaining balance of trade helps to maximize trade volume over the long run.

6. Related Work

The work reported in this paper fits within the general stream of work on agent-based computational economics (ACE), defined by Tesfatsion [15] as "the computational study of economics modeled as evolving systems of autonomous interacting agents". Work on ACE can be categorized into descriptive models, focusing on the constructive explanation of emergent global behavior and normative models, focusing on mechanism design. Our work fits in the latter category since we are using our simulator primarily to evaluate the effectiveness of various brokering strategies.

Mizuta and Steiglitz [1] present an agent-based simulator for dynamic online auctions in order to study some of the dynamic interaction that is not easily captured in the usual theoretical models. They focus on modeling two types of bidding agents: early bidders who bid at any time during the auction period and snipers who wait till the last moments to bid. Agent behavior is specified in terms of a number of parameters, such as the watch probability, limit value, and valuation function. The authors describe bidding behavior they have observed in some online auctions and show that their simulation captures this behavior. They list comparing the simulation results with real-world data as one particularly important area for future research.

Mizuta and Yamagata [2] describe their Artificial Society and Interacting Agents (ASIA) simulator. The simulator is designed to provide a general framework for building agent-based economic and social simulations. The ASIA system provides facilities for agents and users to create agents, dispose agents, and send messages through a MessageManager. The system leaves the concrete design of agent hierarchy, social structure, and individual agent behavior to the user. The authors describe the application of their simulator to three problems: an asset market, a dynamic online auction, and international greenhouse gas emissions trading.

Yamamoto and Sycara [4] present a new scheme for buyer coalition formation and use a simulator to evaluate its effectiveness against a traditional group buying scheme and an optimal allocation scheme. They simulate a reverse auction in which buyers post asks and sellers make bids for items with volume discount prices. Their simulator contains numerous parameters, such as price decreasing ratio, number of buyers, ratio of buyers preferring multiple items, and each buyer's reservation price.

Bohte, et al [3] present a competitive market-based mechanism for allocating banner advertising space. Each supplier bids in a single-bid sealed auction for the banner spaces to be presented to particular consumers. They use simulation to evaluate the effectiveness of their approach. The simulation consists of suppliers and consumers, which are distinct. Each supplier's goal is to maximize immediate profits, with bidding strategies represented as piece-wise linear functions that are learned using an evolutionary algorithm. Consumers are characterized by three different types of purchasing behavior. They show agreement of their simulation results with traditional economic theory of efficient markets.

Table 2. Simulation with optimization (O) and without optimization (NoO)

Case	Average absolute balance per week ($)	Decrease in absolute balance O vs NoO (%)	Average Total Trade volume per week ($)	Increase in Trade volume per week O vs NoO (%)
No Balance Opt.	1,874,709.66		69,121.43	
Balance Opt.	1,421,266.44	24%	97,037.04	40%

7. Conclusions

We have presented a simulator for commercial barter trade, an area that has been little studied up to now. The simulator is built using a data-driven approach by learning probabilistic models of company purchase behavior from transaction history data. We quantitatively evaluated the accuracy of the simulation by comparing simulated trade with actual trade along four different dimensions. Results show a high degree of agreement between the simulated and actual trade.

A number of features can yet be added to our simulator to make it more realistic. These include the tendency of businesses to stick with suppliers and the effect of a company's current balance on its tendency to spend its trade dollars.

Acknowledgements

The brokering module was built by Namthip Rujikeadkumjorn. We thank Bob Bagga and Chris Haddawy at BizXchange for their time and patience in helping us to understand the operations of trade exchanges and for generously providing the transaction data. We thank Hugin Expert for providing us the use of the Hugin Researcher software which was used to learn and evaluate the purchase prediction models.

References

[1] H. Mizuta and K. Steiglitz, Agent-Based Simulation of Dynamic Online Auctions. In *Proceedings of the 2000 Winter Simulation Conference*, pages 1772 – 1777, 2000.

[2] H. Mizuta and Y. Yamagata, Agent-Based Simulation and Greenhouse Gas Emission Trading. In *Proc. Of the 2001 Winter Simulation Conference*, pages 535 – 540, 2001.

[3] S. Bohte, E. Gerding, H. La Poutre, Competitive Market-based Allocation of Consumer Attention Space. In *Proc. EC'01*, pages 202 – 204, Tampa, Florida, Oct 2001.

[4] J. Yamamoto and K. Sycara, A Stable and Efficient Buyer Coalition Formation Scheme for E-Marketplaces. In *Proc. AGENTS'01*, page 576-582, Montreal, Canada, June 2001.

[5] F. Alkemade, H. La Poutre, H. Amman, Intermediaries in an Electronic Trade Network. In *Proc. EC'03*, pages 200 – 201, San Diego, California, June 2003.

[6] International Reciprocal Trade Association, 2001. www.irta.net.

[7] International Reciprocal Trade Association, 1999. http://www.irta.net/barterstatistics.html.

[8] B. Cresti. Welfare improving barter in imperfect competition. Discussion Paper N. 2002/71, Center for Operations Research and Econometrics (CORE), 2002.

[9] B. Cresti. U.S. domestic barter: An empirical investigation. Discussion Paper N. 2003-5, Institut de Recherches Economiques et Sociales (IRES), 2003. http://www.ires.ucl.ac.be/DP/IRES DP/2003-5.pdf.

[10] Breese, J. S., Heckerman, D., Kadie, C., "Empirical Analysis of Predictive Algorithms for Collaborative Filtering", *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, Madison*, WI, July 1998.

[11] C.E. Metz, Basic principles of ROC analysis. *Sem Nucl Med* 8: 283-298, 1978.

[12] Law, A. M., Kelton, W. D., *Simulation Modeling and Analysis*, McGraw-Hill, 2000.

[13] Nash, J. E., Sutclifee, J. V. River flow Forecasting through Conceptual Models, *Journal of Hydrology, vol. 10*, pages 282-290, 1970.

[14] P. Haddawy, C. Cheng, N. Rujikeadkumjorn, K. Dhananaiyapergse, Balanced Matching of Buyers and Sellers in E-Marketplaces: The Barter Trade Exchange Model, *Proeedings of the Sixth International Conference on Electronic Commerce,* Delft, Netherlands, Oct 2004.

[15] L. Tesfatsion. Agnet-Based Computational Economics: Growing Economies from the Bottom Up, *Artificial Life*, vol. 8, no. 1, pages 55-82, 2002.