Proceedings of The National Conference
on Undergraduate Research (NCUR) 2003
University of Utah, Salt Lake City, Utah
March 13 – 15, 2003

Adaptive Agent-Based Simulations of Global Trade

David Koziol and David Roberts
Department of Mathematics
Colgate University
13 Oak Drive
Hamilton, NY 13346. USA

Faculty Advisor: Dr. Dan Schult

Abstract

To examine the strategies of nations participating in the global economy, an agent-based mathematical model was developed and implemented as a parallel algorithm using communicating processes. Each process represents a nation and contains five sub-processes which model the negotiations between nations. Based on Ricardian comparative advantage theory, a series of equations were derived determining all necessary factors affecting trade partner selection and trade levels. A discrete algorithm for each nation to use was designed utilizing these equations. Trade negotiations were modeled as two player games with offers and acceptances. Each agent was given the goal of maximizing its one step utility increase based on a Cobb-Douglas utility function and current commodity levels. The model was verified by running multiple simulations and varying initial parameters. While the expected relationship between quantity of trade and utility increase did not occur, there was a significant relationship between number of trades and specialization. Nations chose a product that they could produce easily and used their comparative advantage to trade for other goods. By redesigning the process where exchange of goods takes place and including a more detailed production function, we expect to see utility increase correspond to levels of trade.

1. Statement and Overview

Modern international trade theory is based upon a publication by David Ricardo and has endured since 1817. This publication established the law of comparative advantage, which expanded upon the absolute advantage theory proposed by Adam Smith. In Smith's theory, it was asserted that if one nation were more efficient in the production of one good and less efficient in the production of a second commodity, mutually beneficial trade would be made possible by specialization. A nation could specialize in and export the good in which it has an absolute advantage. Ricardo refined this theory and proved that even if one nation is more efficient in the production of both goods than another nation, there could still be mutually beneficial trade. This nation could export the good in which its efficiency advantage was larger and import the good in which its advantage is smaller. The former commodity is referred to as the good in which the nation has a comparative advantage.

Ricardo based his theory upon the labor theory of value, which asserts that the price of a commodity is fully dependent on the amount of labor needed for production. For this implication to hold, it must be true that labor and capital intensities are equivalent across the production of all commodities and all labor must be homogeneous. This is not the case in real markets, so the law of comparative advantage must be established via another basis. One such basis is opportunity cost theory. Using this approach, the cost of a commodity can be stated in terms of the amount of another commodity that must be given up in order to create enough resources to produce an additional unit of the first commodity. It then follows that a nation has a comparative advantage in a good when its opportunity cost to produce the commodity is lower than the second country's.This theory will be the foundation of our mathematical model of international trade. Utilizing the opportunity cost theory, the terms of trade that offer mutually beneficial outcomes can easily be established with the knowledge of the production possibility frontier, which shows the different combinations of two goods that can be produced using all of the available resources (for details, see following section and appendix). For the feasibility of the model, it will be assumed that opportunity costs are constant, implying that the production possibility curves are straight lines. It can further be seen that with trade, both nations can consume beyond their production frontiers, and thus can be seen as establishing a greater utility through specialization and exchange.

Examining the theory associated with trade shows that maintaining the lowest possible trade costs and remaining as close as possible to a laissez-faire economy will have the greatest positive impact on the utility of the nations involved. Through our mathematical mo del, there are several issues that we would like to address, with the most pressing that of the utility levels with and without trade. Theoretically, higher utility levels should be reached with trade than in a market with isolated countries. Also of interest are the production levels of each good. Using theory, it will seem as though a nation will specialize in the good of their comparative advantage, and will produce considerably more of one good than the other, if not only one good. An accurate model should produce results that are consistent with this. The other main area to be examined is the volume of trade of a nation per cycle and the other nations with which it typically trades.

Through experimentation, we can show how these relationships can differ based on the costs associated with trade. It is expected that with lower costs associated with trading, trade volume will be highest and each nation will achieve the highest possible utilities. As trade costs (tariffs, etc.) increase, trade volume is exp ected to decrease to the point that most countries produce in isolation and have lower utilities, similar to those in pre-trade equilibrium. There will be few relationships formed in this scenario. More realistically, trade costs can be high between certain agents and low between others. This would be expected to form relationships (similar to EU or NAFTA) between countries. Utility would be increased from pre-trade, but not as high as in the free trade scenario. Trade volume would also be expected to be somewhere between pre-trade and free trade levels.

2. Model Description

The trade model that we have developed consists of a fixed number of agents, each representing a nation, which exist on a non-spatial domain. In this simplified model we will assume that there are only two commodities, sugar and spice. There are two means by which the agents can obtain these goods: either through domestic production or via trading with other nations. The goal of each individual agent is to increase their wealth and maximize their utility function through a combination of production and trade.

As stated above, agents desire to maximize their utility function via any means possible. All agents have several characteristics. Foremost among them are the production rates of sugar and spice. The production rates of the agents represent the maximum level of production of one commodity given that they dedicate all of their resources to that commodity. Agents also have a level of wealth that represents their current holding level of sugar and spice and a consumption rate (or metabolism rate) for each good that represents the rate at which they diminish their wealth level. Each of the afore-mentioned characteristics is initially exogenously assigned randomly to each agent prior to the beginning of the passage of time in the simulation.

In each generation of the simulation, all agents participate in a cycle determining what goods they will obtain and how to obtain them followed by the procurement of the good and an updating of overall wealth and utility. The first step of the cycle is the determination of what commodity, sugar or spice, should be obtained in the generation. This selection is based on whether sugar or spice will have a greater contribution in maximizing the utility level of each agent. The utility level is a function of sugar and spice, and is defined by (1),

$$U(w,v) = w^{(m/(m+n))} * v^{(n/(m+n))} \qquad (1)$$

where w represents sugar wealth, v represents spice wealth, m is sugar metabolism and n is spice metabolism. The good that will have a greater contribution can be established through the marginal rate of substitution (MRS) derived from the ratio of partial derivatives of an agent's utility function with respect to w to that with respect to v. The MRS determines which good enhances utility more, and this good should be sought for in the current trade cycle.

Each agent will participate in this step of the cycle and determine whether they wish to obtain sugar or spice. The next step is the determination of how to obtain the desired commodity. This step involves communication between agents and the determination of compatible partners for trade. Since this model is non-spatial, all agents are able to communicate with one another, unlike spatial models that involve communication restrictions. The criteria for a trade to take place are as follows: (1) agents must be trying to obtain opposite goods and (2) the agents must have a comparative advantage in the good that they want to export relative to the potential partner. For a mathematical description of how this is determined, see Appendix A, step 2.

Having established the criteria necessary for partners to be compatible a matching mechanism determines which trades will take place. Each agent is allowed to submit an exogenously determined number of offers, called O_f, to traders which are determined to be compatible, first by condition (1) and then by condition (2) above. The offers extended will correspond to the agents that will provide the most beneficial terms of trade. Each agent will also be allowed to accept an exogenously determined number of offers, called A_c, based on which offers are determined to be the most preferential. In our model, O_f will be equal to A_c. Preference will be based on the range of mutually beneficial terms of trade (Appendix A, step 2). The range of mutually beneficial terms of trade will be based upon the internal rates of substitution between sugar and spice of the participant nations. If there is a large range of mutually beneficial terms of trade (i.e. a larger difference between internal rates of substitution), this yields the possibility for greater gains available via trade.

The determination of partners will be deferred so that all trade offers will be made prior to the establishment of any trading relationships. This deferment of partnership choice is done to enable agents to make intelligent trade decisions based on all information made available. Every agent will examine all offers made from every other agent before making a trade decision. Trades will then take place when there are reciprocated offers between agents. At the end of each round, agents will produce both goods regardless of trade status.

Once an agent makes a trade decision, the next step for those with established relationships will be completing the trades. After the trade partnerships are established, terms of trade and volume of trade must be determined. The terms of trade will be modeled by taking the multiplicative average of the internal rates of substitution of sugar for spice of each nation in the trade (Appendix A, step 4). This manner of averaging should maintain a reasonable terms of trade between the two nations that is contained in the set of mutually beneficial terms. This process will be carried out for each of the trading partnerships.

With the terms of trade established, the final step needed for the trade to take place is the determination of the volume of trade that will occur between the two partners. The method chosen will determine the trade volume based upon the current wealth of the agents, their production levels, and the number of possible accepted offers per generation (see Appendix A, step 4. option 1). This method has drawbacks as discussed in section 4.

Given the overall volume of trade that has taken place, the agents will consume based upon their metabolism rates and the local agent variables will be updated to reflect the occurrences of the previous generation. Each agent's sugar and spice wealth levels will be updated according to (2) and (3):

$$w_{t+1} = w_t - m + (\text{quantity obtained via trade/production}) \quad (2)$$
$$v_{t+1} = v_t - n + (\text{quantity obtained via trade/production}). \quad (3)$$

The final step in the trade cycle will be gathering information and updating aggregate "report" variables that will function as the overall measure of market performance. The model will calculate and output the average utility of the agents, average sugar wealth, average spice wealth. Also output are total volume of trade occurring in the generation and the average exchange rate. These outputs will be used in experimentation to demonstrate the utility and wealth effects that are associated with higher volumes of trade.

3. Results

3.1. implementation

Due to the nature of the autonomous agents, each stepping though time in unison making trade decisions, we decided to use Java with CSP. Java with CSP, or jcsp, is a set of libraries for the java programming language. CSP stands for communicating sequential processes. The sequential nature of many programming languages (which lack the ability for processes to communicate concurrently) makes it difficult to accurately simulate the parallel characteristics of real world situations. Since CSProcesses are capable of communicating as if they were run in parallel, we are better able to mimic the workings of a real world trade market. Due to the parallel nature of simulations in general, jcsp is an ideal choice for implementing any model.

With the help of the jcsp library, we created a series of autonomous individual agents, represented as CSProcesses. These agents were given random initial conditions according to the constraints we imposed. Once the initial conditions were set, each agent became responsible for making it's own trade decisions. Through the use of communication links, the agents were able to decide for which good to trade, which agents could provide that good, and which agents could provide the most mutually beneficial trade. Each agent reported its specific activities as well as contributing that information to overall aggregates. Thus, we are easily able to monitor any changes in aggregate and individual behavior that occurred as the result of the manipulation of initial conditions.

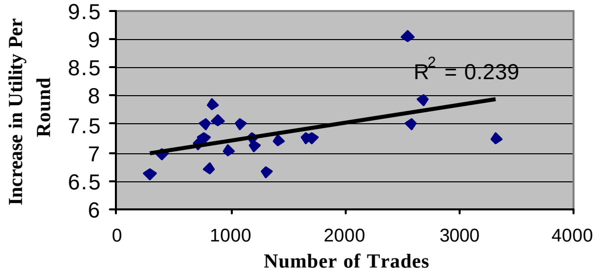

Plot 1 Average per round increase in utility vs. number of trades

3.2. experimentation

We defined the initial "base case" of our experiment to include 20 agents, each of which has randomly assigned initial characteristics. These characteristics were determined via a pseudorandom number generator and took on the following initial values: production levels from 5 to 15, wealth levels from 4 to 14 and metabolism levels from 1 to 5. Our initial level of trade costs was set equal to 0 to symbolize a world with free trade. With these randomly assigned initial characteristics, the agents stepped through 100 iterations of the model as described in section 2 and appendix A. They went through 100 rounds of production, consumption, trade, their characteristics were updated and an output was generated after each round. The aggregate levels were then noted at the end of the 100 iterations.

The randomization of the initial conditions required many trials to generate aggregate outputs that we could compare to those generated upon the variation of initial parameters. Therefore, the base case of the model was run for 20 trials with the outputs after the 100 iterations noted for each trial. The purpose of this base case is two-fold: to provide a platform from which the variations can be measured and to determine the initial relationship between trade and utility levels. The design of the model was such that increases in the levels of trade should be directly related to increases in the utility levels of individual agents and the aggregate average utility levels of all the agents.

The base case yielded results that were somewhat consistent with the initial expectations of the model, although the relationships were not as strong as expected. A scatter-plot of the number of trades that took place in each trial and the average aggregate utility increase per iteration showed that higher utility levels were associated with higher levels of trade (see Plot 1). The trials that had higher overall trade levels typically resulted in higher average utility levels further indicating that trade was beneficial to the individual agents.

3.3. variation

The first initial condition we varied was trade cost. We hoped to uncover the detrimental effects that tariffs have on the free trade market on the aggregate level. We anticipated seeing, on the individual level, agents with relatively lower production rates being excluded from trades. This would translate into a lower overall number of trades taking place on the aggregate level. Initially we raised the cost from 0 to 1 and reran the simulation 15 times. In general, we found that higher trade costs did indeed produce slightly lower levels of utility, but the change was not as significant as we had initially anticipated. What was significant was that increasing the trade costs eliminated relatively less beneficial trades, resulting in a lower number of overall trades taking place in each round (See plot 2).

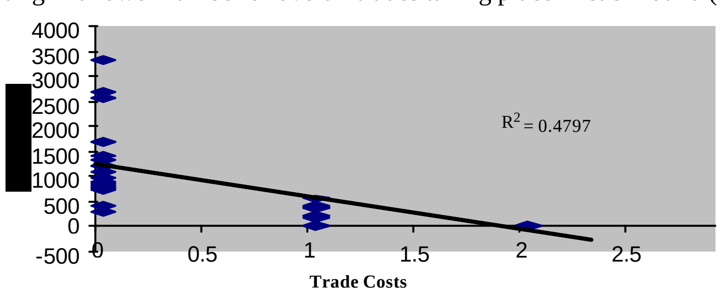

Plot 2 Number of trades vs. trade costs

Next, we attempted to isolate the effects of the range of production levels initially assigned to each agent and held constant in all prior trials. To do this, we returned trade cost to zero and increased the production rate range from 10 to 15. With a larger variation in the production levels, we expected to see more mutually beneficial trades completed each round and thus contributing more to aggregate utility. We expected this because trade is determined to be mutually beneficial based on the internal rate of substitution for each participating agent that is based on production levels. Similarly, since wealth is not just a function of trades, but of production level as well, we expected to see an even greater increase in the aggregate utility level. To test this hypothesis, we again ran trials at three different trade costs, 0, 1, and 2. The first set of 15 trials at cost 0 revealed that with more variation in the production levels, there would be more trades taking place. As a result, overall utility was seen to significantly increase. Our second set of trials, with trade cost set equal to 1, yielded similar findings.

Perhaps the most notable effect of increased production variation occurred when we conducted our trials with trade cost equal to 2. Contrary to our findings in the base case, four out of ten trials yielded trades. This is undoubtedly due to the larger range of production levels increasing the difference in internal rates of substitution, thus making trades more beneficial. As seen before, with other trade costs, the rate of increase for utility per round is significantly larger than that of our base case.

Having seen reasonable results from the modification of production level range, we decided to hold this constant again at 10 and modify base production value. We increased this from 5 to 10 and conducted 10 trials with all variables set back at defaults except for base production level. Since the variance of the production levels has not changed, we didn't anticipate seeing the number of trades increase or decrease. This is due to the fact that the variance in the internal rates of substitution (used to calculate beneficial trades) will remain similar to that of the trials with lower base values. However, it was likely that utility would increase at a greater rate due to the increased wealth based on higher overall production levels. This is exactly what we found. Regardless of trade cost, utility level increased at a greater pace compared to the base case. While the number of trades was comparable with that of the base case, once a trade cost was introduced, no trades took place at all. This is likely due to the fact that a good can be acquired at less cost internally.

Since both increased baseline production and increased production range had significant effects on aggregate output, we decided to probe further by varying both at the same time. Surprisingly, no distinct pattern of trade increase occurs. However, as we might have suspected, aggregate utility reached unprecedented levels.

Since our previous attempts to manipulate the output had been so successful, although we did not expect to see any significant changes in data, we decided to try manipulating another variable: metabolism. Metabolism is used in calculating utility, marginal rate of substitution, and is used to update wealth at the end of each round. As a result of this, we expected only a slight drop in utility and in fact our test results illustrated near exact numbers of trades with only slight drops in aggregate utility.

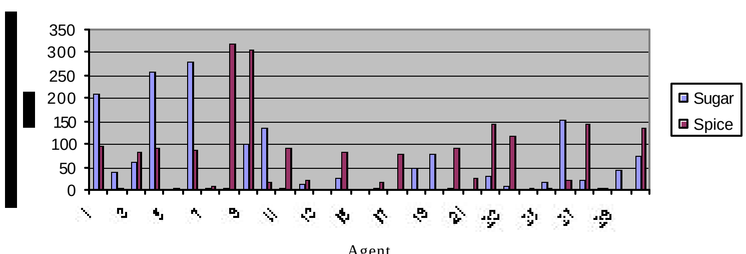

Chart 1 Specialization

Perhaps the most exciting experiment that we ran was testing for specialization within individual agents' trade patterns. We structured the experiment in the same manner as the base case, holding the arbitrarily assigned parameters constant. We then generated 100 iterations of the model and examined the characteristics of individual agents that developed via this simulation. By randomly selecting 10 of the 20 agents and outputting their frequency of sugar and spice trades, we were able to examine the amount of sugar and spice traded for by each agent during each round. This process was repeated three times and the output was charted. The results of these trials provided significant evidence of trade specialization in our model (See Chart 1). Many cases revealed nearly complete specialization while others resulted in partial specialization. A few showed roughly equal number of trades for sugar and spice. This indicates that agents are specializing in the goods that they are relatively more efficient in producing. Occasionally complete trade specialization would occur, with the agents only trading for one good and obtaining the other good via production alone.Allowing for trade costs and larger variations in the production levels of the agents did not eliminate this result, although the extent of its occurrence differed. Allowing for more variation held the amount of specialization essentially constant, while increasing the trade costs caused less specialization due to the lower number of trades taking place. In this case, several agents can be seen as essentially being isolated in the model and not completing any trades. The agents that did trade often specialized in the same manner as was viewed in the previous trials.

3.4. robustness of the model and sensitivity to initial parameters

To examine the robustness of our model, we completed similar experiments to those above while examining variations in the initial arbitrarily assigned parameters. We examined the variations in the model that would result from varying the initial number of agents in the model and the maximum number of trades that each agent can complete in each round of the experiment. While this can be expected to change some numbers, we were hopeful that it would not have a major impact on the results.

Consider the number of agents in the simulation. Our initial testing was done with 20 and we expected to be able to repeat our results with differing levels of agents. We tested this by varying the number of agents to 10 and 30. At each of these levels we conducted base case trials to verify that our results could be reproduced. We examined the effect of changing trade costs and the specialization that resulted at these differing agent levels. By decreasing the number of agents to 10, the results we saw were very similar to those in the base case. An increase in trade costs to 1 from 0 decreased the volume of trade that occurred, with many of the trials yielding no trades at all. This is very similar to the base case results. Furthermore, the change in number of trades seems to vary directly with variations in the number of agents. This variation also produced similar results when considering specialization. Agents were seen as specializing in the good that they are relatively more efficient in producing

When we increased the number of agents to 30, we began to see slightly differing results with regard to trade costs. While increasing trade costs, we saw that many agents continued to makes trades despite the higher cost. This indicates that the greater the number of agents, the more likely it is that pairs of agents resistant to slight increases in trade costs will exist. This is consistent with the overall increase in number of trades that take place when increasing the number of agents. Once again, specialization remained present despite the increase in the number of agents. With 30 agents, we viewed complete specialization with some agents while others allocated their trades rather equally between the two goods, indicating that these agents are among the most efficient in producing both sugar and spice. This condition will be more likely when there are more agents in the simulation.

4. Model Critique

Our model is by nature suggestive and not predictive. It presents an oversimplification of the world by limiting the number of goods in the market to two. For a trade model to be predictive, it would have to allow for more goods to be traded and more complex trading structures. The trading structure in our model only allows for direct exchange between two goods, yet in the real world trade primarily takes place through the exchange of currency for goods. The international trade market is too complex to accurately model with a simulation such as ours. There are too many variables that must be considered. One must look at exchange rates, established trading relationships (such as free trade areas), varying tariffs, varying transportation costs as a result of the spatial component, and the nearly infinite amount of market goods. There are also trades that take place, and are blocked, for political, rather than economic reasons. By simplifying this model to only include the basic theoretical principles of international economics, we can begin to see some of the relationships that are observed in the real world.

Our model is able to depict several of the relationships that we had hoped to see. Most important in terms of the economic theory behind trade is the evolution of a large degree of specialization. By utilizing internal production efficiency rather than marginal rate of substitution to establish the reciprocation and rate of exchange between trading partners, this should be the case. Individual agents should specialize in the goods that they are most efficient in producing. This will especially be the case when their efficiency ratio is highly skewed in favor of one good. In this case, the agents will specialize in one good and trade for the other in any given round. This can be seen in our model and is the best result that we have established.

Another factor that demonstrated results similar to those viewed in the real world was our implementation of trade costs. By increasing the costs associated with completing trades, the number of trades decreased significantly. This is what happens in the real world as the tariffs associated with trade are typically implemented to protect the interests of domestic producers and lower the number of traded goods that enter the market. Despite having the same basic effect on trade as real world tariffs, our execution of costs was an over-simplification. The trade costs as we invoked them do not change the actual exchange rates, but instead eliminate the less beneficial trades. While this assumption may be valid in some limited circumstances, there are many cases in which the trade costs will have additional effects. Further, had we envisioned that minute changes in trade cost would have such a large effect we may have implemented it as a real number rather than an integer, enabling us to test small changes in trade costs rather than whole integer steps.

The most disappointing result from our model was the lack of a large effect of trade on aggregate utility levels. While in the base case we saw a slight correlation between the number of trades taking place and an increased utility level, this relationship did not hold for the majority of the trials that we ran. In the real world, trade occurs with the goal of making all participants better off and the theory that we implemented in our model had this goal in mind. By structuring trade decisions on marginal rate of substitution, all agents were acting to increase their utility at the greatest rate possible. Also, using internal efficiency levels in the selection of the fittest partners for trade should have caused wealth and utility to increase at a faster rate.There are many areas in which our model can be improved to better incorporate behavior that will maximize utility. One clear improvement involves the complete implementation of production specialization. This cannot solve the problem without other improvements also being made. For example, trade volume determination was oversimplified and caused agents to continue trading beyond or halt trading at a point that is not beneficial. We implemented this strategy (see appendix A, step 4, option 1) for simplicity; however, in retrospect it goes against the intent of our model and perhaps our alternative strategy would have produced better results.

Another flaw we uncovered during testing is the potential for negative wealth. We realize that production rate must be larger than metabolism rate. For those agents who do not often make trades, a metabolism rate that exceeds production will cause wealth to drop below zero. While this in itself is not a logistical problem in our model, it can cause one we will discuss later. Negative wealth does not have a meaning and does not exist in the real world. It is possible to be in debt of a good), but in the real world one cannot consume more goods than they have. The logistical error in our model is caused by zero sugar wealth. While calculating MRS, agents must divide their spice wealth by their sugar wealth. If sugar wealth is zero, which is possible if it may become negative, calculating MRS will generate an error. This is a slight problem, as a zero sugar wealth should in theory be allowed. The result of this, having to maintain production level above metabolism rate, reduces the need for trade. If agents were to have metabolism rates greater than their production rates, they would have a greater need for trade to increase their utility levels. Our model does not allow for this.

Because of the need to initialize the model with production levels far greater than the metabolism levels, the overall wealth and utility of the agents will increase at a very fast rate beyond realistic rates. While wealth levels may increase marginally over time, they will typically not exceed a certain amount. When the stock of a good reaches a certain level, it is no longer in the agent's best interest to continue to increase the wealth in this good because it will not be able to consume the excess. This rapid rate of growth may also be partially due to the simplification of our "world" in which there are only two goods being exchanged. By maintaining metabolism and production rates similar in size, this problem would be resolved, but as we said, our model cannot account for this.

In conclusion, while our model may have a deficiency in its identification of the utility effects associated with trade, we are happy with the specialization that was present in our simulation. This is an area that most other trade models are unable to account for. By utilizing the Ricardian theory of comparative advantage rather than a simple MRS to determine partners for mutually beneficial trade, agents with high production levels in one good tend to specialize in that good and trade for the opposite. The simple MRS utilized in other models is in constant flux and therefore does not establish a good of choice to specialize in. The implementation of a more in depth measure of utility including production rates in addition to the already present metabolism and wealth levels should allow for a model of trade that will account for both specialization and effect of trade on utility.

5 Appendix A

Aggregate Parameters:

N - total number of agents

O_f – maximum number of offers to me made

T – cost of trading (i.e. tariff level)

A_c – maximum number of offers to be accepted

With A_c being set to equal O_f in the model.

Aggregate Variables:

U – average utility

W_{total} – average sugar wealth

V_{total} – average spice wealth

E – total volume of trade

R – average exchange rate

U_{std}.= standard deviation of utility

W_{std}.= standard deviation of sugar wealth

V_{std}.= standard deviation of spice wealth

Agent Parameters:

P_w – sugar production level

P_v – spice production level

m – sugar metabolism rate

n – spice metabolism rate

Agent Variables:

w – sugar wealth v – spice wealth U(w,v) = w^(m/(m+n))*v^(n/(m+n)) (utility function)

$w_{t+1} = w_t - m$ + (quantity obtained via trade/production)

$v_{t+1} = v_t - n$ + (quantity obtained via trade/production)

The quantity obtained will be positive for imports and negative for exports.

Algorithm for Stepping Through Time:

Step 1: Commodity Determination

Each agent will calculate MRS:

MRS = (the partial of U with respect to W)/(the partial of U with respect to V)

If MRS > 1, agent will try to obtain sugar If MRS < 1, agent will try to obtain spice

Step 2: Communication (Offer Determination)

If MRS > 1, for each agent with MRS < 1 determine if there exists a comparative advantage for agent *i* in sugar and the extent of this comparative advantage.

If $P_w^i/P_v^i > P_w^j/P_v^j$ then there exists the necessary comparative advantage. If this holds, calculate the difference between the two values. If the difference is larger than T, then offers will be made to the O_f agents with the relatively highest comparative advantage levels. This process is symmetric for MRS < 1.

Step 3: Acceptance of Offers:

Option 1: Check for any reciprocated offers. An agent can accept at most A_c reciprocated offers. If one or more exists, use option one below to obtain a maximum volume and communicate that to the other agent(s).

Option 2: Instead of determining reciprocated offers, after potential trade partners are located, randomly establish up to A_c relationships for each agent.

Step 4: Determination of Terms and Trade Volume:

To determine terms of trade, take multiplicative average of P_w^i/P_v^i and P_w^j/P_v^j call it W_t/V_t. (This is the square root of the product of the internal rates of production for each agent. This will give us the ratio of sugar to spice in the exchange.

Option1:

Maximum volume of sugar exchanged during an individual trade is determined by $(w_i + P_w^{\,i} - m_i)/A_c$.

Maximum volume of spice exchanged during an individual trade is determined by $(v_j + P_v^{\,j} - m_j)/A_c$.

The trade amount will be determined by the minimum of the maximum sugar volume and maximum spice volume multiplied by the exchange rate.

Option 2:

W_t/V_t (defined differently for each relationship) sugar units are traded for one spice unit with each partner. Alternate trade between partners, trading in this manner. When MRS = 1 for a given agent, this agent will stop trading because additional trade will not increase its utility. Each agent can continue trading with other partners until each partner reaches this state and thus eliminates all possible trading relationships.

Step 5: Aggregate Updates

After all trades have taken place for a given round, each agent will consume sugar and spice based upon their given metabolism rates, m and n, and produce sugar and spice based on their given production rates P_w and P_v. A reporting and updating of the agent variables and aggregate measures of market status will follow this.

Go to step 1 with updated agent and aggregate variables.

Note: The specifications above are based upon an agent trying to obtain sugar. They will be symmetric if an agent is trying to obtain spice.

6. References

Bearce, David H. and Eric O'N. Fisher. "Economic Geography, Trade and War." Ohio State University, Nov. 2001.

McFadzean, David and Leigh Tefatsion. "A C++ Platform for the Evolution of Trade Networks." ISU Economic Report Series No. 39, Iowa State University, Ames, June 1999.

Salvatore, Dominick. International Economics, 7th Edition. New York, John Wiley & Sons: 2001.

Tefatsion, Leigh. "A Trade Network Game with Endogenous Partner Selection." ISU Economic Report Series No. 36, Iowa State University, Ames, April 1995.

Tefatsion, Leigh. "How Economists Can Get Alife." Iowa State University, March 1997.