

An Agent-Based Model of Trade with Distance-Based Transaction Cost

Kumar Venkat
Surya Technologies, Inc.
4888 NW Bethany Blvd., Suite K5, #191
Portland, OR 97229
kvenkat@suryatech.com

Wayne W. Wakeland
Systems Science Ph.D. Program
Portland State University
Portland, OR 97207

Keywords: Agent-based modeling, agent-based computational economics, trade, transaction cost, environmental cost.

Abstract

This paper describes an application of agent-based modeling to investigate the effect of a distance-based transaction cost on trade. Long-distance trade is rapidly increasing, but may ultimately be constrained by our ability to move material goods between sellers and buyers. Unlike information exchange, trade in material goods is dependent on the price of oil and vulnerable to future scarcities of oil. In addition, there are growing concerns about greenhouse gas emissions from long-distance transportation. Our purpose in this study is to take the first step in understanding the impact of a distance constraint on free global trade using a simple artificial economy. We use the perspective of agent-based computational economics to model two different scenarios of random initial allocations of goods among traders, and investigate the response of the economy as a distance-based transaction cost is applied. We show that a geographically skewed initial allocation of goods performs poorly, while a more uniform initial distribution responds in a highly resilient way as the transaction cost is varied. Underlying this resilience is the emergence of a stable trade network that has some of the properties of scale-free networks.

INTRODUCTION

The rapid expansion of global trade is outpacing our understanding of the social and environmental consequences of trade. Trade is conducted increasingly over long distances and relies heavily on fossil fuels to transport goods from producers to consumers [1]. Even perishable products such as fruits and vegetables travel an average of 1500 miles (as of

a few years ago) within the United States. When imported foods are added to the mix, the average distance from farm to the dinner table increases significantly. Most non-perishable items that we purchase – from toys to clothes to computers – are typically manufactured thousands of miles away.

As long-distance trade increases, we are also witnessing significant increases in the price of oil, and growing concerns about greenhouse gas emissions from transportation and other sources. How would global trade change if some of the environmental costs – such as fossil fuel depletion and greenhouse gas emissions – are internalized? Alternately, what would be the effect of substantial increases in fuel costs due to increasing demand and decreasing supply?

Concurrently, information technology has become commonplace in every part of the world. Information is increasingly available at a low cost in real time, with possibly a much lower environmental cost than long-distance transportation of goods. Historically, one of the primary obstacles to free global trade has been the cost of accurate and timely information, including the ability to easily find and negotiate with trading partners. Electronic trade exchanges and other Internet-based mechanisms are now beginning to remove this obstacle. But the ultimate constraint to trade may well be our ability to physically move material goods from sellers to buyers over long distances at an acceptable cost including externalities. This constraint might lead to a restructuring of the fast-growing society of global traders, and stimulate new kinds of trade relationships and networks.

Our purpose in this study is to take the first step in understanding the impact of a distance constraint on free global trade using a simple artificial economy. We quantify this constraint as a distance-based transaction cost that traders must pay in order to transport goods to each other. We are interested in

understanding how such a constraint would affect the outcomes of trade, such as the utility or welfare extracted by traders and the prices that traders pay for goods. We are also interested in exploring any geographical patterns that might emerge as a result of the transaction cost, such as trade networks with specific topologies.

We proceed further by building on the substantial recent research in agent-based computational economics or ACE [2], which studies economies as evolving systems of autonomous interacting agents. ACE relies on computational laboratories to study the evolution of decentralized market economies under controlled experimental conditions. It is particularly appropriate for modeling the profit or utility seeking behavior of individual traders who are at different geographical locations with different sets of neighbors and have varying amounts of goods to trade.

METHOD

We formulate the trade problem as follows, adapting a simple barter economy used by Wilhite [3, 4]. Our artificial world consists of 1024 traders spaced uniformly in the four quadrants of a rectangular space, as shown in Figure 1. Each trader is an agent who remains at a fixed location, and is able to trade with others who may be at other arbitrary locations. Traders are presumed to find potential trade partners and negotiate the terms of trade through mechanisms that are independent of their locations, such as globally-accessible electronic trade exchanges.

Each trader starts out with an initial endowment of two durable goods, g_1 and g_2 , ranging from 0 to 1500 units each. The two goods suffer no degradation over time and serve as assets that can be exchanged. There is no production and the aggregate stock of goods changes only to account for the transaction cost as described later.

The initial allocation can follow two distinct scenarios, maintaining nearly equal amounts of g_1 and g_2 in our artificial world:

- “Globally mixed random” (GMR): There are no regional differences. Each trader gets random quantities of the two goods such that the total quantity of both goods together is exactly 1500 units.
- “Local comparative advantage random” (LCAR): The eastern half of the world has more g_1 than g_2 , and the western half has more g_2 than g_1 . Each trader in the east

receives at least 1200 units of g_1 and no more than 300 units of g_2 . Each trader in the west receives at least 1200 units of g_2 and no more than 300 units of g_1 . The actual amounts are allocated randomly such that each trader starts with a total quantity of 1500 units.




Figure 1. Location of traders in the artificial world.

Each trader attempts to maximize the same symmetric Cobb-Douglas utility function, $U = g_1 * g_2$. Traders are rational, non-strategic and myopic such that they do not plan for future opportunities or mislead potential trade partners.

Trade proceeds as follows in this artificial world:

- In each round of trade, traders are chosen in random order and each trader is given a chance to initiate four consecutive trades.
- The trader then searches the world and finds the best possible trade partners for the four trades.
- Two traders consummate a trade if their marginal rates of substitutions are different, and if the utilities ($U = g_1 * g_2$) of both traders increase as a result.
- Trade price between agent i and agent k is determined by the following rule: Price = $(g_{2i} + g_{2k}) / (g_{1i} + g_{1k})$.

- In each trade, the initiating trader buys or sells one unit of g1 in exchange for an appropriate quantity of g2.
- Trade rounds proceed as above and terminate when there are no further profitable trading opportunities.

In addition, traders take into account a transaction cost based on the distance between the two traders. This transaction cost reflects some degree of internalization of the real environmental costs of long-distance trade, including fossil fuel depletion and greenhouse gas emissions. This cost can also simply reflect a high cost of petroleum-based fuels, quite apart from environmental costs.

In each trade, the cost is paid by traders in proportion to the quantity of g1 or g2 bought and the distance that the goods must be transported. The cost is subtracted from the quantity of goods received by each trader in a trade. Traders evaluate this cost in advance and proceed with a trade only if it would still increase their utilities.

The cost is computed as follows: Transaction cost = distance * quantity of goods bought * unit distance cost. For the purposes of this study, the unit distance cost can be one of the following: 0, 0.001, 0.05, 0.1, 0.15, 0.2, or 0.25. Here, 0.001 would be a very low cost, while 0.25 would be very high.

We assume that other trade-related activities, such as finding trade partners and negotiating the terms of trade, incur negligible transaction costs in relation to the distance-based cost. This could happen, for example, if these other activities are carried out over the Internet and largely automated through appropriate software.

We have used NetLogo [5] to develop an agent-based model of the simple barter economy described above. The purpose of the model and the subsequent simulation experiments is to generate new understanding and insights about the effects of a distance-based transaction cost.

Our study of the simulation results starts with a close examination of the model economy in terms of macroeconomic behavior. We then investigate other aspects of the artificial world, specifically the emergence of trade networks. We characterize the structure of these networks and extract information from the simulations to explain their origin.

Simulation results for such simplified trade scenarios cannot be validated directly against real-world data. Therefore, we have used the following

verification criteria to ensure that the model is internally consistent and accurately reflects our barter economy.

As in a real economy, repeated interactions between individual traders should give rise to macroeconomic patterns, and these emergent phenomena should be recognizable. They include metrics such as trade price, trade distance, and traders' final utilities, as well as details of trade such as the number of trades and the number of searches. As the unit distance cost is varied, these metrics should change in ways that should be readily explainable.

In addition, the results should not qualitatively change when critical model parameters, such as the total population of traders and the maximum initial allocation of goods, are varied within reasonable limits. These parameter variations, including variations of the unit distance cost, constitute a detailed model sensitivity analysis.

We report results for a single population of traders and a fixed maximum initial allocation due to space constraints. Each data point in the following figures was obtained by averaging the results from three replications of a particular simulation using three different pseudo-random seeds.

RESULTS

We present our main simulation results in this section. Figures 2 through 9 display the salient features of trade in this barter economy as a function of the unit distance cost.

Figure 2 shows that the average trade price for the GMR ("globally mixed random") allocation is stable and very close to 1 as expected, since there are nearly equal amounts of g1 and g2 in the world. The standard deviation is low and stable as well. The distance-based transaction cost has almost no effect because traders can easily find trade partners nearby who have different initial endowments.

In contrast, the average trade price for the LCAR ("local comparative advantage random") allocation is highly sensitive to transaction cost. As the cost increases, traders have to contend with a shrinking pool of potential trade partners who may not have the right ratios of g1 and g2 in the LCAR scenario. Traders who are situated close to the border between the eastern and western regions might be able to find attractive deals across the border, but in general, local scarcities make the trade price highly volatile and location-specific.

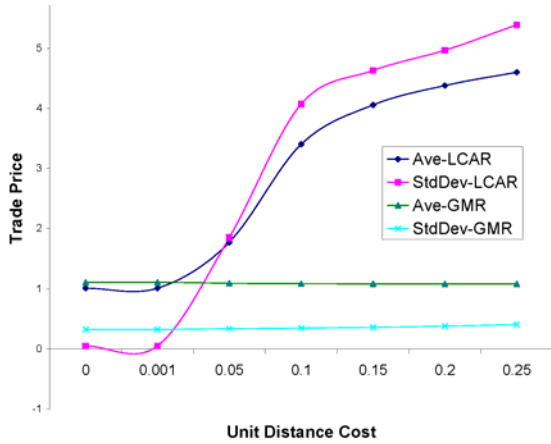


Figure 2. Trade price versus transaction cost.

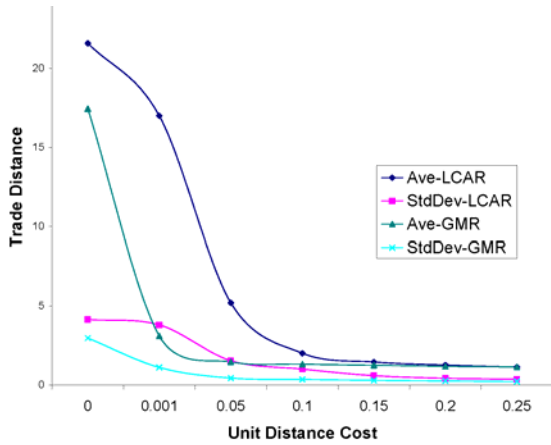


Figure 3. Trade distance versus transaction cost.

Figure 3 shows that the average trade distance declines rapidly as the transaction cost is increased. The trade distance is higher initially for LCAR, since traders have to look farther for profitable trade opportunities. But a unit distance cost of 0.05 or larger makes long-distance trade unprofitable in general and reduces the trade distance. Above this point, the total number of trades declines in the LCAR scenario as seen in Figure 4, while it is nearly unchanged for GMR due to the proximity of good trading partners.

Figure 5 shows that increasing search effort is required to find good trade partners as the unit distance cost increases. In the case of GMR, there is a modest but steady increase in the number of searches. For LCAR, the number of searches increases quickly at first but declines after the cost

exceeds about 0.05 (coinciding with the decline in the number of trades), suggesting that there are significantly reduced trade opportunities at higher transaction costs when the initial allocation of goods is geographically skewed.

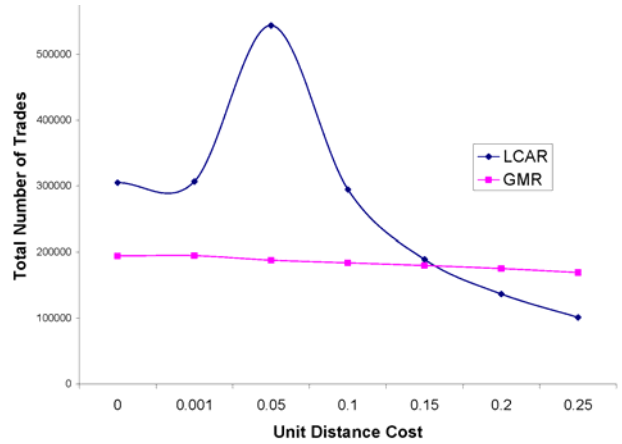


Figure 4. Total number of trades versus transaction cost.

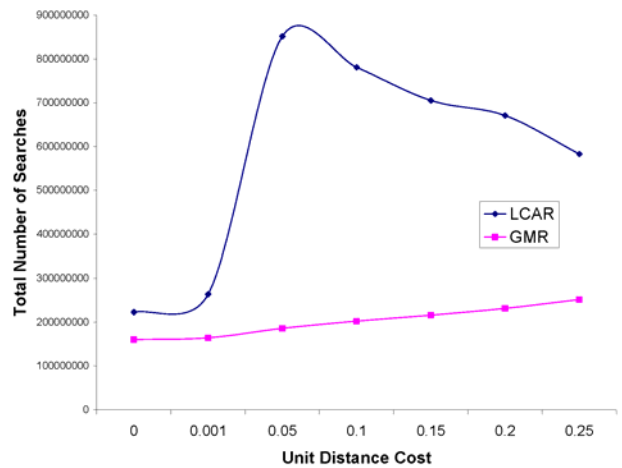


Figure 5. Total number of searches versus transaction cost.

Consequently, the final average utility of the trader population decreases modestly as the unit distance cost is increased in the case of GMR, but undergoes a large drop at about a cost of 0.05 for LCAR, as seen in Figure 6.

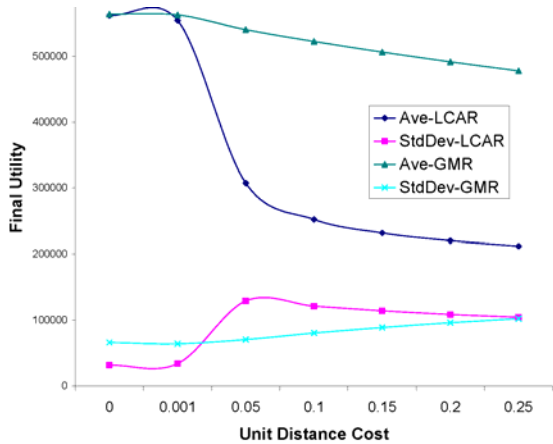


Figure 6. Final utility versus transaction cost.

The final utility for LCAR is about the same as the GMR scenario when there is no transaction cost, which demonstrates that unconstrained trade can efficiently move goods between traders and achieve a level of utility that is nearly independent of initial allocations. LCAR does require more searches and trades in order to overcome geographical differences in the initial allocation.

As the unit distance cost increases in the LCAR scenario, geographical differences in the initial allocation translate into lower final utilities, more volatile prices, more dispersion in the final utilities, and significantly more effort to achieve a given level of average utility. Our model at this point is producing global economic patterns that are consistent and rational.

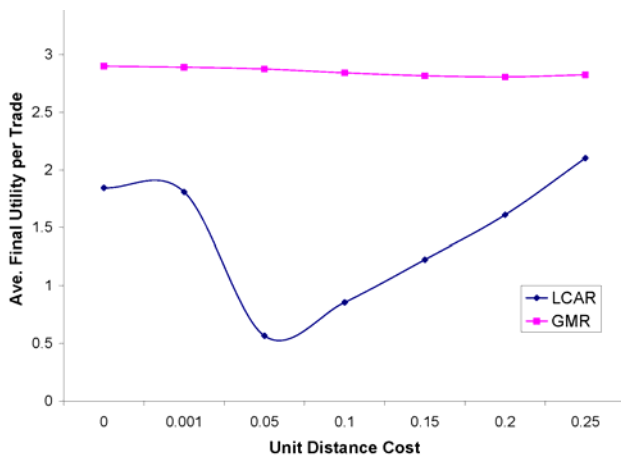


Figure 7. Trade efficiency versus transaction cost.




Figure 8. Search efficiency versus transaction cost.

Given that the final utility of the trader population is an important economic outcome in our trade model, Figures 7 and 8 show the efficiency with which this outcome is achieved as the unit distance cost is varied. In the GMR case, the final average utility per trade is nearly constant regardless of the transaction cost, while the final utility per search decreases modestly due to the slightly increased search effort as a result of the transaction cost. The trade and search efficiencies in the LCAR case are much worse, particularly as the unit distance cost is initially increased, and then the efficiencies improve at higher costs.

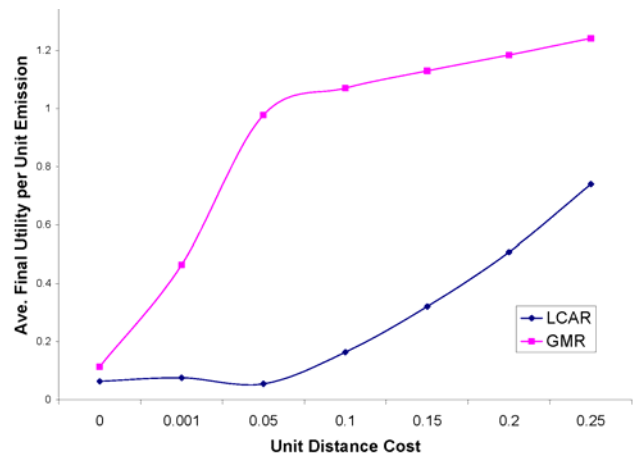


Figure 9. Environmental efficiency versus transaction cost.

Figure 9 shows the environmental efficiency of the economy, as measured by the final utility per unit of greenhouse gas emissions, in response to the transaction cost. We assume that greenhouse gas

emissions produced by each trade are proportional to the quantity of goods and the shipping distance. The environmental efficiency in the GMR case increases rapidly at low costs and then stabilizes at higher costs. Once again, the efficiency in the LCAR case significantly lags the GMR case, but shows an improving trend at higher costs.

Having examined the functioning of the economy under the distance-based transaction cost, we now turn our attention to the emergence of trade networks. Individual traders are the nodes or vertices in these networks, and transactions between traders form the links or edges.

The GMR scenario for initial allocation of goods has proven to be highly resilient in response to the transaction cost. We investigate the reasons by examining the network topologies for this case. Figure 10 shows the structure of the trade network without any transaction cost during the first round of trade. The dots indicate nodes that have not participated in trades yet. After many rounds of trade, it develops the appearance of a random network. In Figure 11, we see a very different trade network emerging after six rounds of trade under a unit distance cost of 0.05.

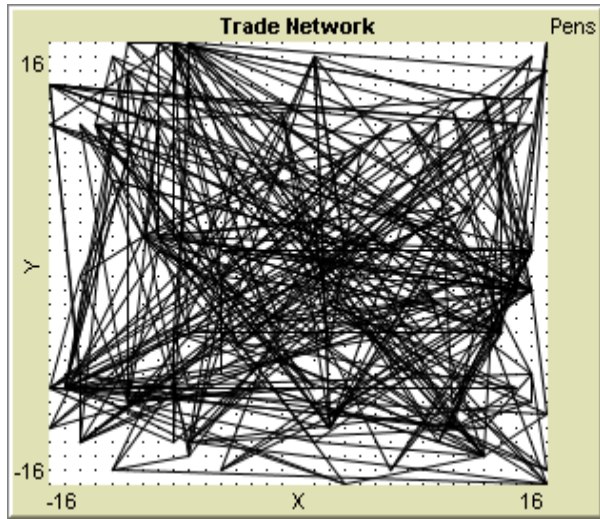


Figure 10. Trade network without transaction cost (GMR).

Figure 12 highlights the difference between the networks depicted in Figures 10 and 11. As the cost increases, the average number of links per node (the average degree of the network) drops quickly. The network no longer appears to be a random network. It starts to form a number of local hubs. Most interestingly, our simulations show that the basic

network structure gels in the first few rounds of trade, and then the structure is reinforced and remains stable during subsequent trade rounds. Figure 12 also shows that the average degree stabilizes to a large extent in the GMR case for unit distance costs at and above 0.05, which correlates with stability and resilience in economic performance. In the LCAR case, the degree changes significantly until a cost of 0.1, and then continues to change at a slower rate for higher costs.

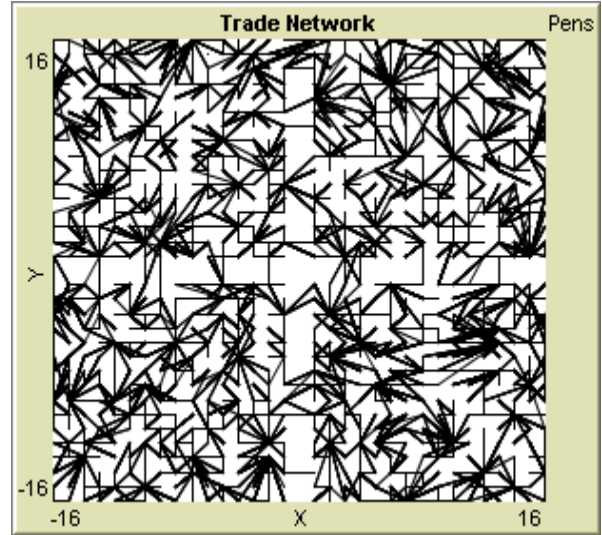


Figure 11. Trade network with Cost=0.05 (GMR).

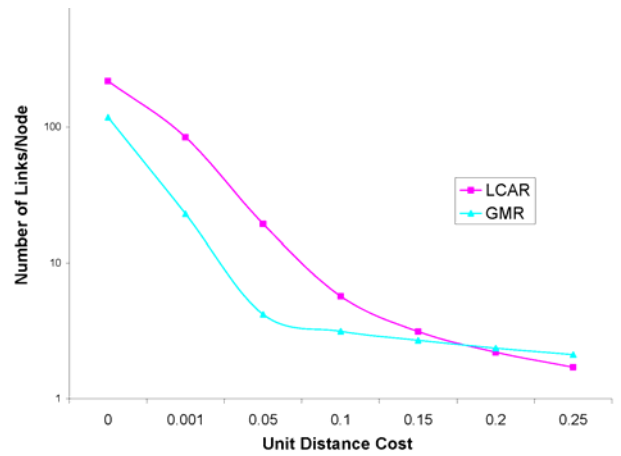


Figure 12. Average degree versus transaction cost.

The network degree distribution in the GMR case for all the unit distance costs is shown in Figures 13 and 14. Without a distance cost, the distribution is somewhat bimodal, with the second peak indicating that a significant number of nodes are of high degree.

When a very low unit distance cost (0.001) is applied, the distribution shifts dramatically to the left where the degree is much smaller.

For medium and high costs (0.05 through 0.25), Figure 14 indicates that most nodes are of very small degree (typically less than 4 links per node), but a few nodes have a larger number of connections. These higher-degree nodes become the hubs in the trade network. The resulting networks resemble power or scale-free networks [6] but the characteristic is limited by the small network size. Over 70 percent of the nodes have at least two connections, suggesting some degree of clustering in the neighborhoods around the hubs.

For comparison, Figure 15 shows that the network in the LCAR case remains random until the cost approaches 0.15, and then it too exhibits scale-free properties. The network topology and the economic performance both stabilize to some extent at these higher unit distance costs.

Our results are consistent with Wilhite's study [4] of fixed economic networks, which showed that the power network is a highly efficient structure that can quickly redistribute goods without suffering from undue search costs and produce a more equitable outcome than many other topologies. Even though we did not start with a fixed network, the emerged network in the GMR scenario displays a similar degree of efficiency for a wide range of transaction costs.

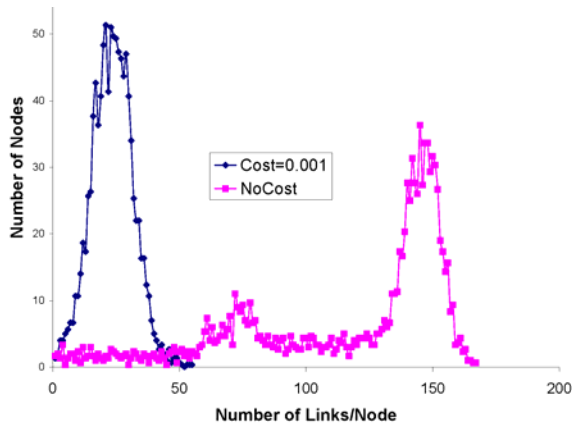


Figure 13. Degree distribution (no cost and very low cost – GMR).

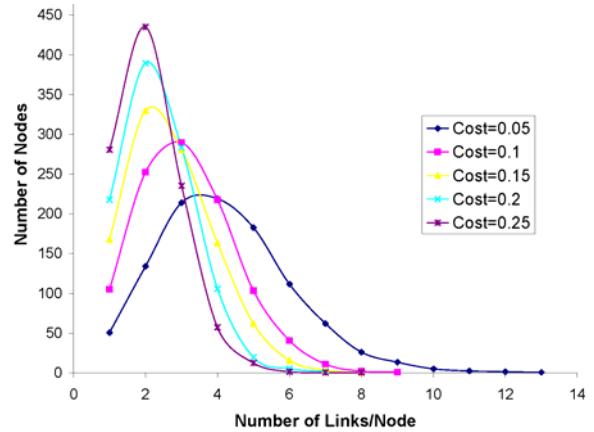


Figure 14. Degree distribution (medium and high costs – GMR).

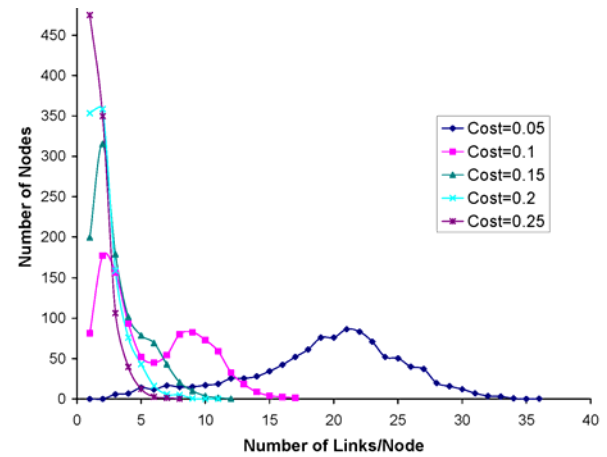


Figure 15. Degree distribution (medium and high costs – LCAR).

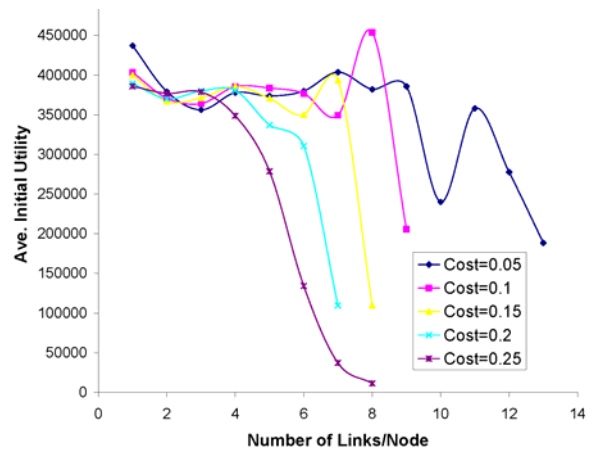


Figure 16. Distribution of average initial utility (GMR).

Figure 16, which plots the average initial utility (initial value of g_1 * initial value of g_2) as a function of

node degree, provides some insight into what causes a trader to become a hub in distance-constrained trade. The nodes (traders) that end up with the largest number of links generally start with low initial utility, and hence the largest motivation to engage in trade in order to improve their individual utility. This motivation to trade makes them local trading hubs for their neighboring nodes. Some of the traders start with low initial utility because their randomly allocated initial quantities of g_1 and g_2 are very unequal.

DISCUSSION AND CONCLUSION

Our simulation experiments show that a distance-based transaction cost has a significant impact on trade in our artificial world. We have two important results in response to the transaction cost. One key result is the profound difference in aggregate behavior of this simple economy in terms of macroeconomic parameters when the method of randomly assigning the initial allocations of goods is changed. The other key result is the emergence of trade networks within the economy.

Without a transaction cost, trade is highly efficient in moving goods between sellers and buyers. A geographically skewed initial allocation of goods in the LCAR scenario produces nearly the same outcomes as the more uniform (but random) initial allocation in the GMR scenario. Average trade price is very close to 1, reflecting the equal amounts of the two goods in the artificial world, regardless of how the goods are distributed initially.

As the transaction cost increases, the GMR scenario is affected only to a moderate degree. Price as well as dispersion in price are nearly the same as in the no cost case. Even though the average trade distance drops to less than 10 percent compared to the no cost case, traders are largely able to compensate by searching their local neighborhoods more, and engaging in more trades. In the worst case (with unit distance cost at 0.25), traders achieve nearly 85 percent of the final utility compared to the no cost case.

The LCAR scenario produces dramatically different results in comparison to GMR. Price becomes highly location-specific, resulting in large changes in the average price and dispersion as transaction cost increases. Trade distance starts out higher than in the GMR scenario because traders have to look farther to find potential trade partners with sufficiently different endowments of the two goods, but the distance drops

to the GMR level at higher costs. Traders try to compensate by engaging in far more searches and trades in their neighborhoods as cost increases initially. But then trading opportunities shrink so much at higher costs that traders ultimately engage in fewer trades than the no cost case. As a result, the final utility drops to less than 40 percent compared to the no cost case.

The resilience of the GMR scenario suggests an evolved structure of trade relationships under the surface. This takes the form of a stable trade network with some of the characteristics of power or scale-free networks. Traders with low initial utility (due to too much of one good and too little of the other) are highly motivated to trade and become local hubs that others can connect to in each small region.

The hub structure and clustering that emerge in our experiments (for GMR, and for LCAR at high costs) are ideally suited for local trade with a small long-distance component. Once such a network has formed, the final utility or welfare of the traders depends critically on each local region having a diverse allocation of tradable goods, as in the GMR case.

REFERENCES

- [1] Venkat, K. 2003. "Global Trade and Climate Change." http://www.greenbiz.com/news/columns_third.cfm?NewsID=26147. *GreenBiz*, December.
- [2] Tesfatsion, L. 2006. "Agent-Based Computational Economics: A Constructive Approach to Economic Theory." *Handbook of Computational Economics*, Vol. 2, Chap. 16. North-Holland, Burlington, MA.
- [3] Wilhite, A. 2001. "Bilateral Trade and 'Small-World' Networks." *Computational Economics*, 18: 49-64.
- [4] Wilhite, A. 2006. "Economic Activity on Fixed Networks." *Handbook of Computational Economics*, Vol. 2, Chap. 20. North-Holland, Burlington, MA.
- [5] Wilensky, U. 1999. *NetLogo*. <http://ccl.northwestern.edu/netlogo/>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.
- [6] Barabasi, A., Z. Dezsó, E. Ravasz, S-H. Yook, and Z. Oltvai. 2004. "Scale-Free and Hierarchical Structures in Complex Networks." *Sitges Proceedings on Complex Networks*.