Adaptive Learning in Evolving Task Allocation Networks[1]

Tomas Klos [a] Bart Nooteboom [b]

[a] *Delft University of Technology*
[b] *Tilburg University*

We study task allocation in multi-agent systems. Task allocation has become a major research topic over the past years [3]. Here, we are particularly interested in allocation of tasks among firms on industrial, interfirm markets. These are traditionally studied using transaction cost economics (TCE). However, as has been widely acknowledged, TCE does not include dynamics of learning, trust and adaptation in its analytical framework, see, e.g. [2].

In our paper, therefore, we extend TCE with notions of trust and loyalty, dynamics and learning, in what we call Agent-based Computational Transaction Cost Economics, the application of the Complex Adaptive Systems (CAS) [1] paradigm to transaction cost economics. In order to be able to do so, we formulate our new theory at the level of individual agents, where these concepts live. In particular, we design an agent-based model, and our theory's refutable hypotheses are derived not by deduction as usual, but by implementing and running our model on a computer.

The Model

TCE takes the 'transaction' as its basic unit of analysis, and analyzes which structural forms should be used for organizing such transactions. If activities are thought of as nodes, and transactions as directed edges between nodes (showing how the outputs of certain activities are inputs to others), then TCE is essentially concerned with the partitioning of nodes into subgroups (firms): edges between nodes within the same firm are organized using hierarchical *firm governance* (the 'make' alternative), while edges between nodes in different firms are organized across firm boundaries using *market governance* (the 'buy' alternative).

We model interactions between buyers and suppliers on an industrial market, i.e. a market for a *component* buyers use to produce a final good which they sell on a final goods (consumer) market.[2] The buyers may buy the component from a supplier ('buy') or produce it for themselves ('make'). In our agent-based model, we let these make-or-buy decisions result from individual agents' decision making, embedded in a (Gale-Shapley-type) matching algorithm that operates on buyer- and supplier agents. Agents' preferences for transacting with each other are calculated as *scores* agents assign to each other:

$$\text{score}_{ij} = \text{potential profit}_{ij}^{\alpha_i} \cdot \text{trust}_{ij}^{1-\alpha_i} + \tau_i,$$

where score_{ij} is the score agent i assigns to agent j, $\text{potential profit}_{ij}$ is the profit agent i can *potentially* make in a transaction with agent j, trust_{ij} is agent i's trust in agent j (which we take as agent i's subjective assessment of the probability that potential profit will actually materialize), and τ_i is agent i's loyalty.

The preferences based on these scores form the input for the matching algorithm which is executed in each timestep of the simulation. The buyer also calculates his own score for himself, and considers any supplier unacceptable who scores lower than this. This then may lead to the buyer being matched to himself (and *making*) rather than to a supplier. To incorporate dynamics and learning, finally, we let the agents update their trust in partners as well as their values for α_i and τ_i (loyalty) using a simple reinforcement learning algorithm.

---

[1] Our paper was published in: J. Sichman, K. Decker, C. Sierra and C. Castelfranchi (eds.), *Proceedings AAMAS 2009*, p. 465–472.

[2] We use the terms 'buyer' and 'supplier' for the agents on the industrial market, and the terms 'seller' and 'consumer' for the agents on the final goods market. A buyer on the industrial market is a seller on the final goods market.

Experiments

We performed experiments with 10 suppliers and 30 buyers in the market, and varied final goods market conditions through the level of product differentiation d. The more differentiated the buyers' products on the final goods market are, the more specific to a certain buyer a supplier's production will be, the less opportunities for economies of scale are present, and the more buyers will therefore choose to make rather than buy. Figure 1 shows that our simulations reproduce this phenomenon, which validates our model.

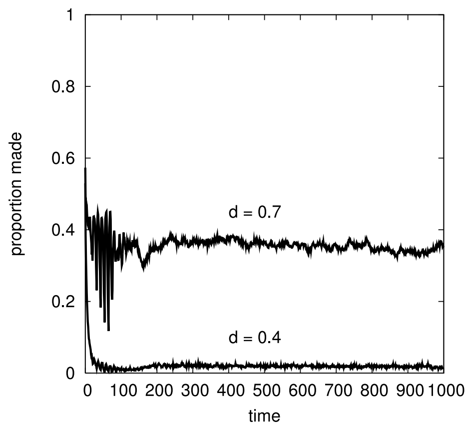

Figure 1: Fraction of production made (not bought) by the buyers.

Of course, the added value of our theory lies in its ability to make statements about the level of individual agents, unlike existing economic theory. In particular, we are able to state hypotheses about, e.g. network formation processes. Figure 2 shows some of these results. Figure 2(a) shows that with low d, most buyers

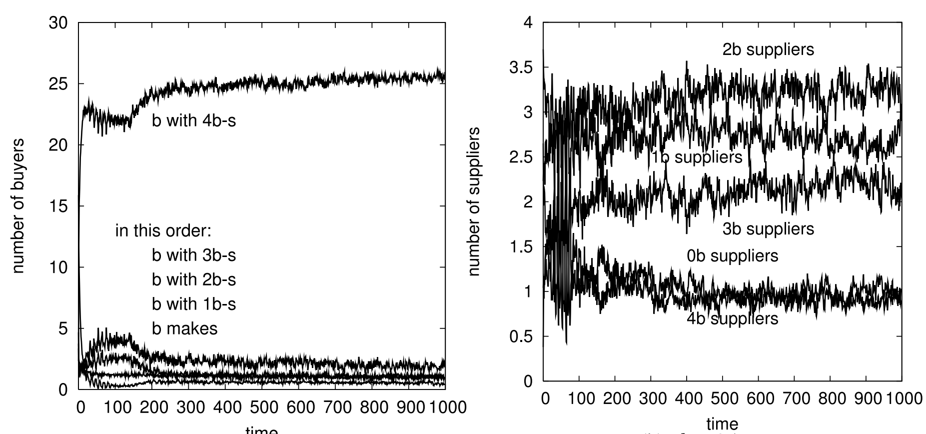

(a) $d = 0.4$.

(b) $d = 0.7$.

Figure 2: The buyers' indegree (2(a)) and the suppliers outdegree (2(b)).

buy from highly efficient suppliers (having 4 buyers each). With high d, Figure 2(a) shows that many suppliers survive who supply to relatively small numbers of buyers. The buyers aren't able to coordinate in groups of 4 at 1 supplier per group.

References

[1] J. H. Holland and J. H. Miller. Artificial adaptive agents in economic theory. *Am. Ec. Rev.*, 81, 1991.

[2] B. Nooteboom. Towards a dynamic theory of transactions. *Journal of Evolutionary Economics*, 2, 1992.

[3] M. De Weerdt, Y. Zhang, and T. Klos. Distributed task allocation in social networks. In *AAMAS*, 2007.