

Money Networks in Kiyotaki-Wright Model

Timothy A. Moran¹ and Antonella Ianni¹

¹The University of Southampton, Southampton SO17 1BJ, UK
tam1g09@soton.ac.uk

Abstract

The benefits of money as a medium of exchange are obvious, but existing models of monetary search fail to account for the separate relationships that agents develop between providers and acceptors of money. In this paper, the Kiyotaki-Wright model of monetary search is reproduced as small population agent-based model, and the random matching rule is replaced by a weighted matching model where the probability of successful trades is adaptively reinforced over time. Results show the endogenous emergence of stable patterns of networked trade transactions.

Introduction

Money and its general acceptance as a medium of exchange lie at the heart of most economic activity. Its use offers a convenient alternative to barter, allowing agents who share a belief in its acceptability to trade indirectly using a monetary good that offers them no direct utility. It also offers a decentralised alternative to personal credit arrangements if the acceptance of the money is widespread.

But agents' use of money within a trading environment is highly repetitive and asymmetric. For instance, employees will be paid regularly by a single employer once a month, and then spread out their spending with several vendors during that time period. Existing models of monetary search that treat agent matching as a random process do not incorporate this element of directed search.

This paper begins by introducing an economic search model of money and its use in experiments with real and artificial agents. The model is then implemented as an agent-based simulation and extended* to allow matching patterns that depart from the random matching assumption of the original formulation of Kiyotaki-Wright. Results document the endogenous emergence of networked patterns of trade that prove to be stable over time.

A Search Model of Money

Kiyotaki & Wright (1989) proposed a probabilistic search and matching model that can support monetary equilibria where useful commodities are valued as media of exchange. The economy consists of three types of agent (I, II and III) who can each hold a single unit of one of three goods (1, 2 and 3).

* Preliminary and incomplete.

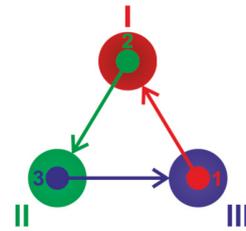


Fig. 1. Production and consumption in the Kiyotaki-Wright model. Type I agents consume good 1 and produce good 2, type II agents consume good 2 and produce good 3, and type III agents consume good 3 and produce good 1

Agents can produce one type of good, but only derive utility by consuming a different type of good. An agent will consume its consumption good immediately, and will produce its production good after consuming. (Thus an agent is never empty-handed.) Since no agent produces its own consumption good, inter-agent trade is necessary for agents to derive utility.

Agents have the opportunity to trade through a random matching process. In every time period, agents are randomly paired and given the opportunity to trade. The model is designed to ensure that there exists no 'double coincidence of wants' (Jevons, 1875) between any two agents. In other words, for trade to take place at least one agent must be willing to accept a good other than its consumption good. (This sets the stage for a good to potentially emerge as a medium of exchange.) Trade only takes place when both agents in a pair value their partner's holding more highly than their own. Thus agents will always accept their own consumption good and they will never trade with an agent holding the same good that they are already holding.

Trade in other goods depends on the trading strategies of agents. To differentiate between the good types, the model imposes different storage costs for each. Letting c_j denote the cost of holding good type j between trading turns, then $c_3 > c_2 > c_1$, meaning that good 3 is the most costly to store and good 1 is the least costly.

Agents attempt to maximise their expected discounted lifetime utility. If they do not believe that any particular good will increase their chance of trading in a subsequent turn then they consider only the physical properties of the goods, and will only accept their consumption good or a commodity that is cheaper to store than their current holding. In this **fundamental equilibrium** type I and type III agents will

never trade directly, as type I agents aim to minimise costs by never accepting good 3 from type II agents. In a sense, type II agents are willing to use good 1 as money, but only because it is cheaper to store than their production good (3).

As Duffy (2001) points out: ‘An agent *speculates* when he accepts a good in trade that is more costly to store than the good he is currently storing with the expectation that this more costly-to-store good will enable him to more quickly trade for the good he desires to consume.’ For a sufficiently high utility of consumption (or, equivalently, sufficiently low storage costs) type I agents are willing to accept good 3 from type II agents, allowing them to subsequently trade directly with type III agents for their consumption good. In this case a **speculative equilibrium** is supported; type I agents are now willing to use good 3 as money, even though it costs more to store than their production good (2).

The trading strategies for each type of agent can be summarised as $a > b > c$, meaning that a is the favourite good and c is the least favourite good. The agent will trade any holding in exchange for good a (the agent’s consumption good), will trade holding b only in exchange for good a , and will trade holding c in exchange for any other good (Fig. 2).

Equilibrium	Type I	Type II	Type III
Fundamental	$1 > 2 > 3$	$2 > 1 > 3$	$3 > 1 > 2$
Speculative	$1 > 3 > 2$	$2 > 1 > 3$	$3 > 1 > 2$

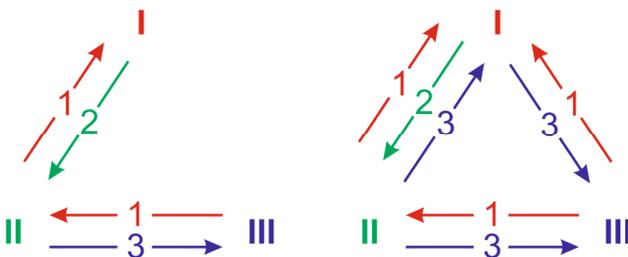


Fig. 2. Trading strategies and resulting trading patterns for the fundamental (left) and speculative (right) equilibria

Extensions to the Search Model

The original model presented only steady-state equilibria in pure strategies. Subsequent work has considered dynamic and mixed-strategy equilibria (Kehoe, 1993), presenting a more generalised model where agents can alternate their play across their two available trading strategies.

The routes by which a monetary equilibrium could become established have been explored using both analytical and agent-based approaches (Alvarez, 2004). Replicator dynamics have been used to demonstrate analytically the dependence of an ultimate monetary equilibrium on initial conditions such as starting strategies, the storage costs of goods, and the proportions of different agent types in the economy (e.g., Luo, 1999; Sethi, 1999, Moran et al. 2013).

The relevance of agent-based approaches to economic modelling is well established (Vriend, 1994; Epstein & Axtell, 1996; Gintis, 1997; Duffy, 2000; Tesfatsion, 2002). Marimon et al. (1990) used classifier systems to allow agents to learn through experience those actions that resulted in positive utility, while Duffy (2001) used experiments with human

subjects to appropriately calibrate an agent-based model. Başı (1999) allowed agents to learn socially through imitation. In general both agent-based and human subject experiments found that social interaction encouraged the use of speculative strategies.

This paper lays the basis for extending the Kiyotaki - Wright model, to allow for the emergence of a networked structure of interaction that is shown to substantially depart from the random matching assumption of the original paper.

Small Population ABM

Real economies consist of finite numbers of participants, with interesting economic behaviour exhibited even in very small economies. Agent-based simulation allows the number of interacting agents to be easily selected. An advantage of running simulations with small populations is that results can be compared to laboratory data from behavioural experiments. Such experiments (Duffy, 2000) have typically used less than 30 agents playing a repeated game for less than 100 periods.

Initialisation

A population size is chosen and an initial population of agents is created, with an equal number of agents of each of the three types. For simplicity the population sizes were chosen to be a multiple of six to ensure an equal distribution across consumption types and to allow all agents to form trading pairs. In the basic model the consumption type also uniquely defines the agent’s trading strategy, with all agents playing fundamental strategies. Agents are initially holding their production goods, representing an economy with no initial endowments or natural resources.

Trade

Each turn agents are randomly paired into potential trading partnerships and attempt to trade according to their pre-defined trading strategies, just as in the Kiyotaki-Wright model. If a successful trade results in an agent holding its consumption good then that agent immediately consumes its holding and gains positive utility by doing so. That agent then immediately produces a new unit of its production good, which becomes its new holding.

At the end of every turn each agent pays the storage cost for its current holding. The utility of consumption (u) and the storage costs for each good (c_1 , c_2 and c_3) are defined globally and are the same for each type of agent. Agents record their lifetime utility. In Moran et al. 2013 the model is expanded to allow agent trading strategies to evolve and this lifetime utility record is used as a measure of the fitness of each agent.

Results

A single run of the simulation consists of the creation of a population of new agents, the interaction of those agents over a number of turns, and data collection to allow the behaviour of those agents to be summarised.

Data was collected for ease of comparison with the results presented in Kiyotaki & Wright (1989). This consisted of the stocks (x) of each good at the end of the turn; the number of transactions (t) involving that good during the turn; the

‘velocity’ (v) of each good; and the ‘acceptability’ (a) of each good. These last two values were chosen as two measures of the ‘moneyness’ of each good, with velocity ($v = t/x$) a more traditional measure (Fisher, 1909) showing the number of transactions weighted by the supply of the good in the economy, while acceptability ($a = t/o$) is the probability that a good will be accepted in trade (Kiyotaki & Wright, 1992), weighting transactions by the number of times a good is offered (o).

Results of a single run are shown for a small population of 90 agents (Fig. 3). Solid lines show the levels at the end of each trading turn. Because fundamental equilibrium trading strategies were imposed the system very quickly settles on the equilibrium levels for stocks, transaction, velocities and acceptabilities, taking less than 10 trading turns to do so.

To record these equilibrium levels, averages are calculated for each good from period 10 onwards, and shown as dotted lines. Even for small populations the results are consistent with large- and infinite-population models.

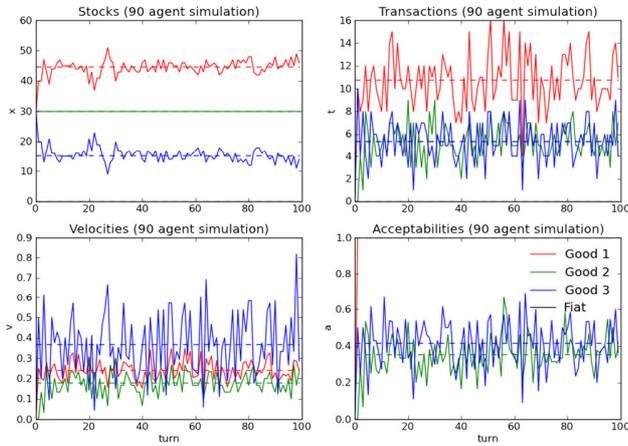


Fig. 3. Results showing stocks, transactions, velocities, and acceptabilities of three goods over time for a single run of the agent-based model with 90 agents

Experience Weighted Money Networks

The existing model assumes a complete trading network, where any two agents may meet and attempt to trade with equal probability. Real trading environments tend to have strong cultural, social or geographical roots, where agents do not meet at random. They instead tend to repeatedly interact with the same small group of the larger population, with many agents receiving money from only a single employer or agency, and repeatedly spending it in the same small selection of shops (Howitt & Clower, 2000).

One explanation for this behaviour is that agents learn to repeat matching behaviours that have benefitted them in the past through a process of reinforcement learning. In this section the random matching process in the Kiyotaki-Wright model is replaced with an agent specific weighted matching process where pairings are more likely if the pair has previously experienced a successful bilateral trade.

Trading Networks in the Agent-Based Model

As before, agents choose a partner and then attempt to trade with that partner. If both the initial agent and the chosen partner prefer each other’s holding to their own, trade occurs.

The difference between the new model and the original Kiyotaki-Wright model is that the choice of partner is no longer completely random, but instead weighted by each agent’s past experience with that partner. Agents that have successfully traded with each other in the past increase their weightings to make it more likely that they will meet again in the future, while agents that have met but were not able to successfully engage in trade decrease their weightings to make it less likely that they will meet again. This approach is a form of reinforcement learning, with the occurrence of trade used to assess the value to an agent of a particular partnership. In the current version of the paper updating of matching probability takes place in an ad hoc manner (detailed below), but we conjecture that in fact any rule that increases the likelihood of repeated trade between agents who have successfully traded in the past would lead to similar results.

The specific partnership updating rules used to produce the results presented in this paper are: in the initial time period, agents weight all other agents equally and so have an equal chance of any particular partnership. This is exactly equivalent to the pure random matching in the original Kiyotaki-Wright model. In this first period the probability of any agent i attempting trade with any agent j is $p_{ij} = 1/(N - 1)$ for $j \neq i$ and $p_{ii} = 0$ (as an agent cannot meet itself).

Trade is successful only if both agents in the partnership prefer their partner’s holding to their own, with preference dictated by the fundamental equilibrium trading strategies from the Kiyotaki-Wright model. In this case the agents swap their holdings, and then consume and produce a new unit of their production good if appropriate, exactly as they did in the Kiyotaki-Wright model.

Agents now update their matching weightings for the particular agent that they have just met. Loosely, if trade is successful the probability of meeting the same partner again is approximately doubled, while if trade is unsuccessful the probability of meeting that partner again is halved.

Specifically, if trade occurs between agent i and its chosen partner k then the numerical weighting of p_{ik} is doubled, but then p_{ik} and all other elements p_{ij} for agent i are re-normalised so that the sum of elements remains 1. This means that the new weightings are: $p'_{ik} = 2p_{ik}/(2p_{ik} + \sum p_{ij})$ and $p'_{ij} = p_{ij}/(2p_{ik} + \sum p_{ij})$, where k is the agent who was chosen as a partner and j is all other agents $j \neq k$.

If trade does not occur, the numerical weighting of p_{ik} is halved and all elements re-normalised. This means the weightings become: $p'_{ik} = 0.5p_{ik}/(0.5p_{ik} + \sum p_{ij})$ and $p'_{ij} = p_{ij}/(0.5p_{ik} + \sum p_{ij})$.

Each agent keeps track of its own potential partnership weightings, making it more likely for agents who have successfully traded in the past to meet again. This particular form of doubling and halving the weightings means that agents are initially likely to experiment with multiple partners, but if they successfully trade with the same partner a few times they develop a very strong trading relationship with that partner and are likely to attempt trade with that partner even after a single failure. Successive failures to trade mean that new partnerships are likely to form.

Directed Networks

In this initial model, updates are performed asymmetrically, with only the initiating agent updating its matching weighting array. This was chosen to allow asymmetric trade networks to arise, where A's preference for trading with B can be independent from B's preference for trading with A. This seems to be a realistic way of modelling monetary trade, where the direction of trade can be important. An alternative would be to allow both partners to update their weightings as a result of any transaction, whether or not they initiated it.

One extension of the model that we plan to pursue is to allow agents to seek out trading partners that are specific to their current holding. For instance, a Type II agent would want to seek Type III agents while holding its production good (3), but would then want to meet Type I agents to exchange the monetary good (1) for its own consumption good (2). This extension is not pursued in these initial results, where agents develop trading patterns that are independent of their current holding.

Results of Weighted Networks

Results are shown for a population of 30 agents. Some representative results from a single run are shown to illustrate how trading networks form.

Nodes represent the individual agents, which are indexed from 1 to 30. Node colour represents each agent's type, with type I's shown in red, type II's in green and type III's in blue.

The left hand panels indicate the realised partnerships that take place in the time step. These partnerships are directed, with an initiator agent choosing their partner based on their current trading partner weightings. An agent can initiate a single trade per time step, but can be the partner in any number of trades initiated by other agents. The direction of the partnership is shown using a rectangular arrow head, with the link pointing from the initiator towards the chosen partner. Colour is used to indicate whether this partnership successfully resulted in trade (green) or not (red).

The right hand panels show the updated trading partner weightings for each agent at the end of the time step, after updating based on the success or failure of the trades shown in the corresponding left hand panel. The width and opacity of each link represents the likelihood of a particular partnership occurring, with faint lines representing unlikely partnerships and strong lines representing very likely or near-certain partnerships. Links are directed from the initiating agent to each potential partner and colour coded by the type of the initiating agent.

Fig. 4 shows three time steps for the partnership reinforcement process. In the first time period (top panel) the partnerships are completely random, as in the original Kiyotaki-Wright model, as agents have not yet updated their partnership probability arrays from their default values. Such random matches are likely to include many cases where agents are meeting other agents of their same type or holding, where trade cannot occur, and the majority of these matches are unsuccessful (shown in red in the left hand panel). The right hand panel shows that agents are effectively indifferent between their potential trading partners by the end of this first period, with all links very weak.

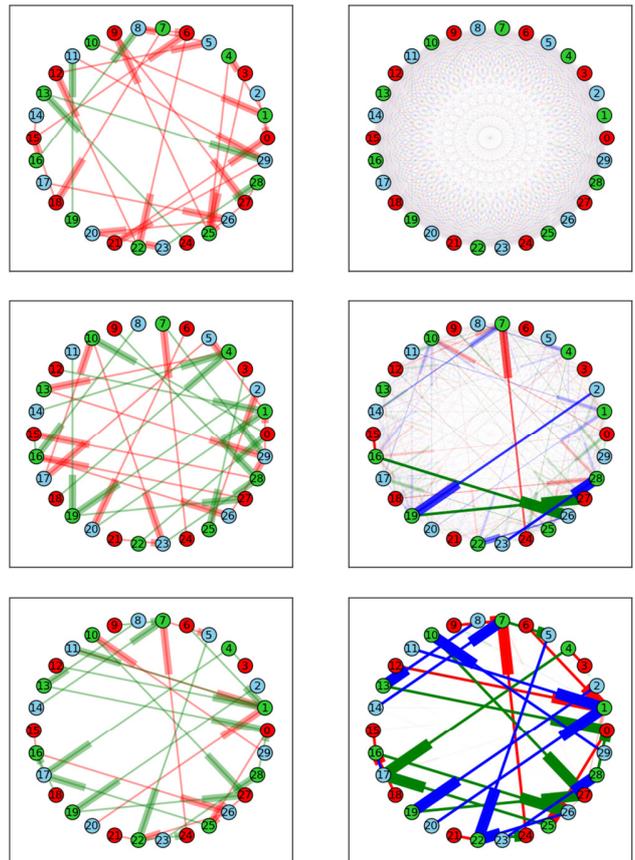


Fig. 4. (left) The actual partnerships that occur in three single timesteps (top: $t = 1$; middle: $t = 50$; bottom: $t = 300$), with successful trades shown in green and unsuccessful trades in red; (right) the resultant weighted trade networks, colour-coded by initiating agent with intensity indicating strength of the probability weighting.

By time period $t = 50$ (middle panel) a small number of agents have formed strong partnerships and most agents have formed weak partnerships with agents who are more likely to offer them successful trade. This becomes very pronounced by time period $t = 300$ (bottom panel) where most agents are almost certain to attempt trade with a unique trading partner. The bottom left hand panel shows that this approach leads to considerably more trade than in the early period or in the basic Kiyotaki-Wright model, with the majority of partnerships in the later periods showing successful trade (green links).

Efficiency of Trade

The proportion of partnerships that do not result in trade are plotted against time period in Fig. 5. These levels correspond to the proportion of red links in the left hand panels of Fig. 4. Starting from the random matching position of around 80% of matches not resulting in trade, agents quickly learn to seek out partnerships that maximise their expectations of trade before reaching a long-run equilibrium where trade is more than three times more likely.

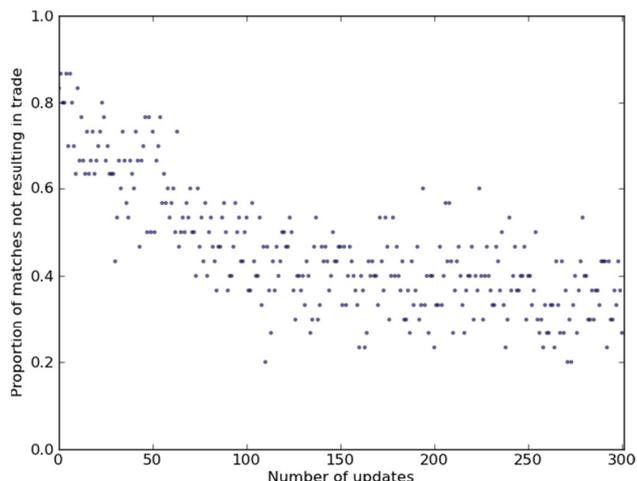


Fig. 5. Partnership selection by agents increases the likelihood of successful trade (results for $N = 30$ agents)

The increase in the level of successful trade relative to the original Kiyotaki-Wright model is clearer when considering a larger population. Fig. 6 compares the proportion of matches that do not result in trade for the new agent weighted matching process (dark blue) with the original random matching process (red). The weighted matching process results in an equilibrium level of trade of just over 60% of matches, compared to 20% for random matches.

The lower probability of a particular match taking place means that the learning process takes longer in the larger population, with the equilibrium level for the matching process now taking more than 500 time steps to reach (compared to less than 100 in Fig. 5). However, the same lower level of just under 40% is reached regardless of population size.

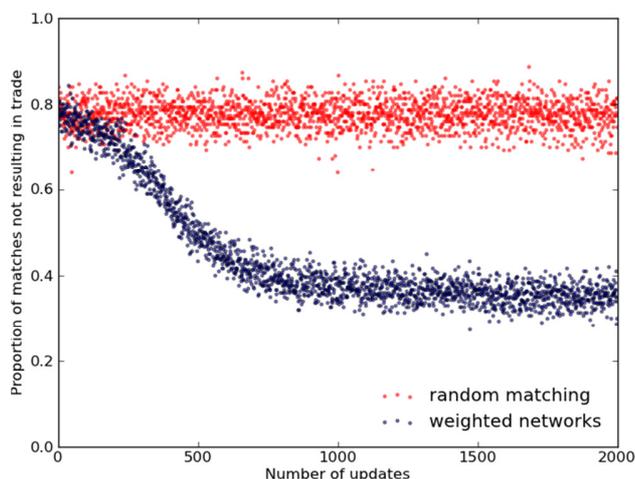


Fig. 6. The network reinforcement model (dark blue) results in successful trade more than twice as frequently as the random matching model (red) (results for $N = 300$ agents)

Discussion

Allowing agents to form trading networks that favour partnerships that have previously resulted in trade results in a higher proportion of trade overall. If agents follow this rule they tend to favour a single partner or a very small group of partners, as tends to be seen in everyday human trading behaviour.

Preliminary results also suggest that this level of trade is even higher if agents are allowed to form trading partnerships that are dependent on their current holding, as the acceptors of the monetary commodity can seek out the most suitable partner depending on whether they are currently holding their production good or the monetary good.

We next intend to include our model of evolving trading strategies (Moran et al., 2013) to see how network formation and trading strategies can co-evolve in a non-equilibrium setting. We will also consider the inclusion of agents who are able to exchange commodity goods for a widely accepted fiat money.

Acknowledgments

This work was supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1).

References

- Alvarez, A. (2004), Learning to choose a commodity-money: Carl Menger's theory of imitation and the search monetary framework, *The European Journal of the History of Economic Thought*, 11:53-78
- Başçı, E. (1999), Learning by imitation, *Journal of Economic Dynamics & Control*, 23:1569-1585
- Duffy, J. and Ochs, J. (1999), Emergence of Money as a Medium of Exchange: An Experimental Study, *The American Economic Review*, 89:847-877
- Duffy, J. (2000), Learning to Speculate: Experiments with Artificial and Real Agents, *Journal of Economic Dynamics and Control*, 25:295-319
- Epstein, J. M. and Axtell, R. (1996), *Growing Artificial Societies: Social Science from the Bottom Up*, The Brookings Institution, Washington D.C.
- Fisher, I. (1909), A Practical Method of Estimating the Velocity of Circulation of Money, *Journal of the Royal Statistical Society*, 72(3):604-618
- Gintis, H. (1997), A Markov Model of Production, Trade and Money: Theory and Artificial Life Simulation, *Computational & Mathematical Organization Theory*, 3:1:19-41
- Howitt, P. and Clower, R. (2000), The emergence of economic organization, *Journal of Economic Behavior and Organization* 41, 55-84
- Jevons, W. S. (1875), *Money and the mechanism of exchange*, D. Appleton and Company, New York
- Kehoe, T. J., Kiyotaki, N. and Wright, R. (1993), More on Money as a Medium of Exchange, *Economic Theory*, 3:297-314
- Kiyotaki, N. and Wright, R. (1989), On Money as a Medium of Exchange, *The Journal of Political Economy*, 97:927-954
- Kiyotaki, N. and Wright, R. (1992), Acceptability, Means of Payment, and Media of Exchange, *Federal Reserve Bank of Minneapolis Quarterly Review*, 16(3):18-20
- Luo, G. Y. (1999), The evolution of money as a medium of exchange, *Journal of Economic Dynamics and Control*, 23:415-458

- Marimon, R., McGrattan, E. and Sargent, T. J. (1990), Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents, *Journal of Economic Dynamics and Control*, 14:329-373
- Moran, T.A., Brede, M., Ianni, A. and Noble, J. (2013), The Origin of Money: An Agent-Based Model
- Sethi, R. (1999), Evolutionary stability and media of exchange, *Journal of Economic Behavior & Organization*, 40:233-254
- Tesfatsion, L. (2002), Agent-Based Computational Economics, Growing Economies from the Bottom Up, *Artificial Life*, 8(1):55-82
- Vriend, N. J. (1994), Artificial Intelligence and Economic Theory, *Many-Agent Simulation and Artificial Life* (1994):31-47