

The Origin of Money: An Agent-Based Model

Timothy A. Moran¹, Markus Brede¹, Antonella Ianni¹ and Jason Noble¹

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 the historical origin of money is less clear. An existing economic model of monetary search is reproduced as an agent-based simulation and an evolutionary algorithm is used to model social learning. This approach captures the way in which different equilibria can arise, including solutions in which one or two goods come to be used as money. In the case where monetary goods have identical properties, multiple equilibria can be reached with a dependence on the starting beliefs of agents. In our analysis we also consider the evolutionary dynamics that allow for a small chance of mutations in strategies. In some cases our findings show evolutionary paths by which use of particular monetary goods can collapse.

Introduction

The economy is a complex adaptive system (Beinhocker, 2007). 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 the value of money as a medium of exchange only arises if that money is widely accepted. The initial growth in the acceptance of money involves the reinforcement of agent beliefs from repeated successful transactions with an emergent form of money, and does not require any centralised coordination. Building an agent-based model of such a system will allow us to assess the plausibility of different historical pathways to the emergence of money, and also to study the conditions that lead to a collapse in the acceptance of a particular monetary system, a topic that economic models have so far neglected.

This paper begins by introducing an economic search model of money and its use in experiments with real and artificial agents. This model is then implemented as an agent-based simulation and extended to allow agents to learn successful trading strategies. Evolutionary paths towards the Nash equilibria are shown.

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).




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).

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. Başçi (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.

Two hypotheses are tested in the following work: an agent-based replication of the Kiyotaki-Wright model is used to test that Kiyotaki & Wright’s results still hold for small populations; and a numerical simulation of trading strategy evolution is used to test the stability of monetary equilibria in the presence of strategy mutations.

Finite Population Model

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. The infinite-population model can be approximated with a large population of several thousand agents, or population sizes less than a hundred can be used to see if the same results hold in small communities. Another advantage of running simulations with small populations is that results can be compared to laboratory data from behavioural experiments. Such experiments 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. When the model is expanded to allow agent trading strategies to evolve, this lifetime utility record will be used as a measure of the fitness of each agent.

Results of the Agent-Based Model

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

Evolving Trading Strategies in the Agent-Based Model

The consistency with the infinite population results indicates that an agent-based model is appropriate for studying the emergence of monetary equilibria. The agent-based model can be modified to allow agents to adapt their behaviour. Instead of imposing unchanging equilibrium trading strategies on the agents, agents are now allowed to adjust their trading strategies based on their relative success. A basic evolutionary algorithm was used to allow successful preferences to be

retained, unsuccessful preferences to be replaced, and new or forgotten preferences to emerge.

Instead of imposing equilibrium trading strategies, agents were given trading strategies that were initially completely random. Regardless of consumption type, agents had a 1/6 probability of being assigned one of the initial trading strategies:

$1 > 2 > 3$	$1 > 3 > 2$	$2 > 1 > 3$
$2 > 3 > 1$	$3 > 1 > 2$	$3 > 2 > 1$

It is worth emphasising that these initial trading strategies are unlikely to be beneficial to the agent, as in many instances they will lead to an agent rejecting its consumption good. However, the evolutionary model will allow agents to adapt their trading strategies to match those that have been successful in the previous generation, allowing us to test whether this model is sufficient for a monetary equilibrium to emerge.

Trading Strategy Fitness

The simulation is broken down into a number of generations (G), each consisting of a number of trading periods (T). At the start of the first generation agents are given random trading strategies as described above. Play within a single generation is the same as described above, with agents being randomly paired and trading when both agents in a pair prefer their partner's holding to their own, given their current trading strategy. Agents keep track of their lifetime utility, which increases whenever they receive and consume their consumption good, and decreases by the storage cost of their holding at the end of each trading turn.

At the end of each generation agents are ranked by their total utility across all the trading turns in that generation. Agents who consume a relatively large amount, or spend fewer turns carrying the goods with the highest storage costs, will have the highest utilities within that generation. This total utility level is used as a measure of the fitness of that agent's trading strategy, with higher fitness trading strategies more likely to survive into subsequent generations.

Imitation and Mutation of Trading Strategies

The agent population is first divided by consumption type. Within each consumption type, the 80% least successful agents are discarded. Each of the most successful 20% of agents then produces four offspring, so that the population size remains unchanged between generations.

Offspring are initially a perfect copy of their parent, with the same consumption type and the same trading strategy. There is then a 10% chance that each child will slightly mutate its trading strategy by swapping the order of two goods in its priority list. The two goods that are swapped in this way are chosen randomly with equal probability of any two goods being selected, i.e. the first and second item may be swapped with $p = 1/20$, or the second and third item may be swapped with $p = 1/20$.

This mutation mechanism means that at most two items are swapped in the trading strategy, with zero chance of a larger mutation or multiple mutations within a generation.

Results of the Evolutionary Model

In all cases a population size of 300 agents (100 of each consumption type) is chosen. Generations consist of $T = 1000$ trading turns, and trading strategies are reproduced and mutated across $G = 50$ generations.

Results of the evolutionary model are plotted as charts that show the dominant trading strategies of each type of agent against generational time. Coloured squares are used to show the proportion of each type of agent using a given trading strategy in a given generation. Colours represent agent consumption type (type I in red, type II in green and type III in blue), with the intensity of that colour showing the proportion of the population who are choosing that trading strategy. Saturated colours represent a trading strategy chosen by all or most agents of a particular type, while very weak colours signify a trading strategy chosen by few or no agents of that type.




Fig. 4. Agent-based evolution of speculative strategy ($u = 100$, $c_1 = 1$, $c_2 = 4$, $c_3 = 9$, average over 20 runs)

When utility of consumption is suitably high, type I agents benefit by adopting the speculative trading strategy, as shown in Fig. 4. Although it takes several generations for agents to adapt, they ultimately settle on the speculative equilibrium.

With lower utilities for consumption ($u < 20$), holding costs become increasingly significant and small populations of agents may not discover the fundamental equilibrium.




Fig. 5. Failure to discover fundamental strategy ($u = 10$, $c_1 = 1$, $c_2 = 4$, $c_3 = 9$, average over 20 runs)

Fig. 5 shows that type II agents prefer a trading strategy that permanently minimises their storage costs by accepting the lowest cost good (1) and never releasing it. Surprisingly, they do not learn to accept their own consumption good. This shows that agents in small-population runs of our model can adopt and persist with trading strategies that were not predicted by analytic approaches. The finding is intriguing but we expect that it would not occur given a larger population and indeed our focus here is not on the discovery of consumption goods (surely a no-brainer in evolutionary terms) but on the origins of monetary exchange.

Evolving Trading Strategies in a Large Population Model

An alternative to simulating the evolution of individual agent behaviours is to simulate the evolution of the population as a whole. In their original paper, Kiyotaki & Wright computed elements p_{ij} of a population array \mathbf{p} , with the elements corresponding to the proportion of type i agents holding good j in the steady state reached after a number of trading steps. In Kiyotaki & Wright these 6 elements of \mathbf{p} were sufficient to completely describe the population because each consumption type held a fixed (fundamental or speculative) trading strategy.

To study the evolution of trading strategies using a population array, new elements need to be added to allow agents of the same consumption type to use different trading strategies. Following a methodology similar to Luo (1999) and Sethi (1999) who studied this problem analytically, the population array \mathbf{p} is reconstructed using 12 elements, with each element of the array representing a triplet of consumption type, holding and trading strategy.

Each consumption type is now permitted two trading strategies, both of which still prioritise that type's consumption good. The interesting question in monetary search is not whether an agent discovers his consumption good (which he must in order to gain any utility), but how that agent treats non-consumption goods in a monetary capacity. In this three-good system the agent can prioritise the remaining two non-consumption goods in two ways: either it can prefer to hold its cheaper non-consumption good (a fundamental trading strategy) or it can prefer to hold its more expensive non-consumption good (a speculative trading strategy).

This is a slightly different labelling than employed in Sethi (1999), which treats consumption types as preferring to hold their production good or their non-production good. In the case of type I and type III agents, their production good is also the cheaper of their two non-consumption goods. But in the case of type II agents, their production good is the more expensive of their non-consumption goods.

Initialising the population

The population array is composed of twelve elements corresponding to one of three consumption types, each of whom may hold one of their two non-consumption goods and one of two trading strategies (fundamental or speculative).

Each consumption type comprises one third of the population, and the population is initialised so that all agents are holding their production good. The proportion of each trading type following each of their two possible trading rules can be varied. As an example, if all consumption types started with equal proportions playing each possible trading strategy, the initial elements of the population array would be:

$$\begin{array}{lll} p_{I(2)123} = 1/6 & p_{III(1)213} = 0 & p_{III(1)312} = 1/6 \\ p_{I(3)123} = 0 & p_{III(3)213} = 1/6 & p_{III(2)312} = 0 \\ p_{I(2)132} = 1/6 & p_{III(1)231} = 0 & p_{III(1)321} = 1/6 \\ p_{I(3)132} = 0 & p_{III(3)231} = 1/6 & p_{III(2)321} = 0 \end{array}$$

where the subscripts represent consumption type-(holding)-trading strategy.

Trading to a steady state in holdings

After initialisation the simulation iterates through an outer loop. At the beginning of the iteration all agents' holdings are reallocated to the production goods of that type. Agents already have trading strategies, either from a previous iteration or from initialisation.

The population shares are updated to reflect repeated matching by agents for the given trading strategies. Any particular match between type i and j will occur with probability $p_i p_j$, with trade occurring if it is beneficial to both members of the pair as in all earlier models.

If as the result of a match an agent ends up holding its consumption good, it immediately consumes it and replaces it with its production good. The population share resulting from such a match is therefore added to the element corresponding to that agent's consumption type, trading strategy and production good. After multiple trading steps a steady-state in goods is reached.

Replication of successful trading strategies

When the holdings reach a steady state – i.e. the holdings on two subsequent turns are sufficiently similar (within a specified tolerance level, set as 10^{-6} for the results in this paper) – the trading phase of the simulation terminates, and the time-discounted expected lifetime utilities of different types of agent are calculated.

Agents are given a great degree of foresight while calculating these expected lifetime utilities. For the reported results, agents were allowed to look-ahead 100 periods using a discount factor of $\beta = 0.9$. Any calculation with more than about 50 periods is a good approximation to an infinitely-lived, perfectly rational agent.

Agents of a given consumption type and holding are then allowed to imitate each other's strategies based on their relative expected utilities. A discretised version of the replicator equation (Weibull, 1995) is used to adjust the population shares across trading strategies for each consumption type-holding pair. The population share for a given consumption type (t), holding (h) and trading strategy (s) is updated as:

$$p_{t(h)s} = p_{t(h)s} \left(1 + r \left(U_{t(h)s} - \frac{p_{t(h)s} U_{t(h)s} + p_{t(h)s'} U_{t(h)s'}}{p_{t(h)s} + p_{t(h)s'}} \right) \right)$$

where s' is the alternative trading strategy for the same consumption type-holding pair, $U_{t(h)s}$ are the expected discounted lifetime utilities already calculated, and r is a scaling factor used to represent selection pressure. The proportion of the population using a particular trading strategy increases if that strategy yields a higher expected utility than the population-weighted average of the two alternative strategies, at a speed proportional to the difference. The proportion playing the less successful strategy will shrink.

After performing a single trading strategy update step, the population is reinitialised to hold their production goods and the next iteration begins with a new round of trading to a steady state in holdings. This process continues until successive updates of the entire outer loop produce no further change in trading strategy share.

After each trading step and trading strategy update the population is re-normalised to ensure that small numerical errors do not result in the creation or destruction of holdings (during the trading steps) or a re-allocation across consumption types (during strategy updates).

The same general results are obtained for less far-sighted agents. The number of future time periods considered in expected utility calculations was chosen to allow relatively rapid convergence to a trading strategy equilibrium, but limiting this amount of foresight only slows the learning process, it does not change the end result. As long as agents consider at least one period into the future, they are able to appreciate the benefits of a monetary good as a medium of exchange.

Results of the Large Population Model

A variety of setups were used to explore conditions under which the different equilibria of the Kiyotaki-Wright model could be reached.

Results are visualised in the trading strategy space of the three agent types. Each consumption type can play one of two strategies: either the fundamental trading strategy that favours holding the lower numbered good (the cheaper good in the conventional setup of $c_1 < c_2 < c_3$), or the alternative speculative trading strategy that favours the higher numbered (more costly) good.

After each trading strategy update, the proportion of each consumption type playing that type's fundamental strategy was recorded, and these proportions used to label the axes of a cube that describes the strategies of all agents in the economy. The 0 of the axis corresponds to all agents playing their speculative trading strategy, while 1 corresponds to all agents playing their fundamental trading strategy.

A selection of starting points was chosen (27 points formed by all possible combinations of [0.25, 0.5 and 0.75] across the three consumption types) and trading strategies allowed to evolve under the rules described above. When plotted these evolving trading strategies tend to trace paths from a uniform three dimensional grid in the centre of the trading strategy space towards one of the equilibria at the corners of the cube. This equilibrium was dependent only on utilities and costs, and not on the particular starting trading strategies.

Results show a representative evolution of trading strategies for the starting mix of trading strategies (1/2, 1/2, 1/2).

Storage costs and utilities have been chosen for consistency with Duffy (2001), with $c_1 = 1$, $c_2 = 4$ and $c_3 = 9$.

Fundamental Equilibrium

With $u = 10$, all paths rapidly converge on the fundamental equilibrium (1,1,1), with all agents preferring lower storage cost goods to higher storage cost goods (type I: $1 > 2 > 3$, type II: $2 > 1 > 3$, type III: $3 > 1 > 2$).




Fig. 6. Fundamental equilibrium (rapid convergence)

Increasing the utility of consumption increases the incentive for type I agents to speculate and experiment with holding a monetary good, as the benefits of more frequent trade are greater relative to the fixed costs of holding goods. With $u = 20$ (and costs unchanged) trading strategies still converge on the fundamental equilibrium, but far more slowly.

Although type II agents still converge rapidly on their fundamental trading strategy, there is very little evolutionary pressure for a mixed population of both fundamental and speculative type I agents to move towards the fundamental trading strategy, as the expected utilities of either trading strategy are very similar.




Fig. 7. Fundamental equilibrium (slow convergence)

Speculative Equilibrium

When the utility of consumption is sufficiently high, type I agents can increase their expected utility by accepting good 3 from type II agents. The additional expense of holding this high storage cost good is offset by the increased expectation of direct trade with type III agents for good 1, the type I agents' consumption good.

Kiyotaki & Wright calculate the critical level at which type I agents should speculate as dependent on the level of utility, the holding costs and the proportions of type II and type III agents who are holding good 1. Type I agents should speculate if $c_{13} - c_{12} < (p_{31} - p_{21})\beta u_1$ (Kiyotaki & Wright, 1989).

This can be seen when the utility of consumption is set sufficiently high, with $u = 100$ in Fig. 8. Convergence to this speculative equilibrium occurs rapidly, with all three consumption types converging on their equilibrium strategy in similar timescales.




Fig. 8. Speculative equilibrium

To produce Fig. 8 an appropriate speed of trading strategy replication (r) needed to be chosen. Lower values of r mean that the system takes longer to reach an equilibrium, but the more gradual replication of trading strategies stops strategies from becoming caught at alternative equilibria.

This occurs because a trading strategy can only be replicated if it still exists within the population. Once entirely eliminated, the replicator equation used above will not allow a trading strategy to re-emerge, as it has a zero weighting in the population average. If r is too high those paths that pass close to the corners (1,1,0) and (1,1,1) may become trapped at these alternative equilibria before reaching the speculative equilibrium (0,1,1).

As well as slowing down the speed of convergence, another way to ensure that the system does not approach a sub-optimal equilibrium due to these numerical errors is to introduce a small mutation rate that allows extinct trading strategies to reappear. In the cases discussed above such a mutation rate will only temporarily move the system away from the equilibrium, but becomes interesting in the case of a mixed equilibria system.

Multiple Equilibria

This framework which had been used to reproduce the evolutionary dynamics described analytically in Sethi (1999) can also be used to consider the case of identical goods, proposed in an example in Luo (1999). If goods are either identical or very similar, the particular type of money used in the economy may have a sensitive dependence on the initial mix of agent trading strategies.

By relaxing the cost ordering of Kiyotaki-Wright and setting the storage cost of goods equal ($c_1 = c_2 = c_3 = 0$), the particular monetary equilibrium depends only on the initial trading strategies used by agents, as shown in Fig. 9 where $(1/3, 2/3, 1/3)$ is a critical point around which trading strategies significantly diverge. Trading strategies started at the 125 points formed by allowing each starting trading strategy to diverge from this critical point by $[-0.02, -0.01, 0, 0.01, 0.02]$.




Fig. 9. Multiple equilibria for identical goods ($u = 100, c_1 = c_2 = c_3 = 0$)

Trading strategies can end at one of the three 'speculative' equilibria, shown in red $(0,0,0)$, green $(0,1,1)$ and blue $(1,1,0)$, or at an equilibrium (shown in grey) where one consumption type continues to be composed of agents playing both 'fundamental' and 'speculative' strategies.

The labels 'fundamental' and 'speculative' are no longer meaningful as all goods have an identical holding cost. However, the particular good that is used as a medium of exchange depends on the equilibrium point that is reached, which depends only on the starting mix of trading strategies:

- $(0,1,1)$: Type I accept good 3, type II accept good 1
- $(0,0,0)$: Type III accept good 2, type I accept good 3
- $(1,1,0)$: Type II accept good 1, type III accept good 2

Strategy Mutation

With trading strategy imitation described by the replicator equation, once agents reach an equilibrium (at a corner or edge of the trading strategy space) they will remain there, as there is no process for forgotten strategies to be rediscovered.

Allowing a small degree of trading strategy mutation after the replication phase allows forgotten trading strategies to return. Each consumption type-holding pair is mutated

independently. In each case a random number is drawn uniformly from the interval $[-0.001, +0.001]$ and multiplied by the proportion of the population playing either trading strategy in this pair. The proportion of agents playing the fundamental trading strategy is then increased by this amount, and the proportion playing the speculative trading strategy decreased by the same amount, with a normalisation step to ensure that this does not result in either proportion becoming less than zero.

In the case of the fundamental and speculative equilibria discussed above these mutations have little effect. Mutations cause the trading strategies to fluctuate around the equilibrium, but strategy replication takes the system back towards it.

However, in the case of the mixed equilibria, strategy mutations can drive the system around the edges of the trading strategy space, permitting sudden transitions from one monetary equilibrium to another. Starting from the equilibrium at $(0,1,1)$, the results of strategy mutations are shown in Fig. 10.




Fig. 10. Trading strategy mutation shifts monetary equilibria




Fig. 11. Cycling through trading strategy equilibria

Fig. 11 shows the speed of these transitions. Trading strategies initially move slowly along the edge of the trading strategy space. Along these edges one consumption type is split into players playing both fundamental and speculative trading strategies, while the other two consumption types are playing a single strategy. The mutation step allows these single-strategy players to experiment with their alternative strategy, which provides an additional incentive for the two-strategy player to shift in favour of the second strategy. Ultimately a critical point is reached where enough of the current two-strategy players are playing the second strategy to cause the rapid transition to a new equilibrium. At this point the two-strategy player has an incentive to play only their second strategy, and a new consumption type begins experimenting with their alternative strategy.

This shows how the gradual shift in the acceptance of one type of money by one type of agent can tip the population structure to the point where an entirely new good becomes accepted as money.

For instance, in the initial move from (0,1,1) towards (0,0,1), with a sudden transition to (0,0,0): Initially good 3 is used as money by type I and good 1 is used as money by type II. Type II agents then increasingly refuse to accept good 1, driving the system towards an equilibrium where their own production good is the unique monetary good. However, at a critical point the system shifts as type III agents begin accepting good 2 as a monetary good.

This process repeats. In each case an agent shifts its trading strategies in favour of creating a monopoly in money production, ultimately resulting in a shift that begins a cycle where the original agents production good is rejected as money entirely.

Discussion

It is encouraging that the models of Kiyotaki & Wright (1989) and Sethi (1999) can be reproduced in both agent-based and numerical simulations that support the original analytic results. The findings presented here explicitly confirm the infinite-population based estimates of Sethi in the limit of very small population size and in the presence of noise in the evolutionary dynamics.

Speculative, fundamental, and mixed equilibria can each be supported if appropriate consumption utilities and storage costs are chosen. If goods are homogenised by setting their storage costs to be equal, multiple equilibria can also be supported; this finding provides a way in to modelling problematic phenomena such as competing currencies (Hayek, 1976) or monetary collapse.

There are many ways in which this framework could be extended. One of the most obvious would be to consider more realistic economies in which more than three types of agents trade more than three types of goods; in which goods differ in their properties such as durability; and in which prices can vary. Another line of extension would be to investigate the effect an evolution that is constrained to real-world social networks has on monetary search. The current 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, suggesting that investigating the

evolution subject to more constrained interaction patterns could be an important step in motivating the maintenance of multiple competing currencies.

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
- Beinhocker, E. D. (2007), *The Origin of Wealth: Evolution, Complexity and the Radical Remaking of Economics*, Random House Business Books, London
- 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
- Hayek, F. A. (1976), *Denationalisation of Money: The Argument Refined*, Institute of Economic Affairs, London
- 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
- 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
- Weibull, J. W. (1995), *Evolutionary Game Theory*, The MIT Press, Cambridge, MA