

A Markov Model of Production, Trade, and Money: Theory and Artificial Life Simulation

Herbert Gintis

SFI WORKING PAPER: 1997-01-006

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu



SANTA FE INSTITUTE

A Markov Model of Production, Trade, and Money: Theory and Artificial Life Simulation

Herbert Gintis
Department of Economics
University of Massachusetts
Amherst, MA 01003

December 28, 1996

Abstract

The paper generalizes the Kiyotaki-Wright trade model by treating the trading period as a finite game, so Nash's theorem can be used to prove the existence of equilibrium, and by treating the economy as a Markov process, so an ergodic theorem can be used to show the existence of equilibria with desirable properties (e.g., in which money exists). A Markov model of trade also allows us to add complexity to the economy without adding corresponding complexity to the analysis of the model's properties. The paper also provides artificial life simulations of the Markov economy suggesting that monetary equilibria are dynamically stable and do not require high levels of learning or information processing on the part of agents.

1 Introduction

The paper analyzes an economy consisting of a finite number of agents, each of whom lives a finite but uncertain number of periods. In each period agents are paired through some stochastic mechanism, and engage in bilateral exchange. Agents produce some goods and consume others, and take in to trade the 'inventory' of goods the agent has acquired through production

and trade, minus the goods the agent has consumed. In each period agents pay a ‘storage cost’ for those goods that they carry in their inventories. We show that under plausible conditions an equilibrium exists in which agents trade with positive probability and one or more goods emerge as media of exchange: goods that are accepted in trade by agents who value them only for their ability to facilitate future trades with other agents.

This model is inspired by (?), which has spawned a fruitful literature.¹ Kiyotaki and Wright’s central insight is that by taking prices as given and ignoring the question of market clearing, the trading process can be modeled in an analytically tractable manner.² We contribute to the model in two ways. First, by treating the trading period as a finite game, we can use Nash’s theorem to prove the existence of equilibrium independent of the specific details of the trading model. Second, by modeling the economy’s a Markov process, we can use the ergodic theorem for Markov processes to prove the existence of equilibria with desirable properties for the economy as a dynamic system. We are thus able to add considerable complexity and structure to the model without adding corresponding complexity to existence proofs.³ While we will in fact deal formally with a relatively simple version of the model, we note where appropriate how the model might be generalized. Ultimately, it is possible that a usable general equilibrium system can be generated, in which both price determination and trade strategies are endogenous, on the basis of a Markov model of production and exchange.

The second contribution of this paper is to provide artificial life simulations of the model suggesting that monetary equilibria are dynamically stable and do not require high levels of learning or information processing on the part of agents. Indeed, our artificial life agents have genomes that are only

¹See, for instance, (?), Kiyotaki and Wright (1991, 1993), Aiyagari and Wallace (1991, 1992), and ?).

²The traditional Walras-Arrow-Debreu model of trade (?) is the opposite of Kiyotaki and Wright’s model, taking prices as determined endogenously by equating supply and demand in all markets, but instead of modeling trade among agents, assuming that a central authority (the so-called ‘auctioneer’) effects all agreed-upon transactions. It would be desirable to develop an analytical model with both trading and endogenous price determination, although it is not obvious whether there are interactions between the two systems that give rise to phenomena not encountered when the systems are treated separately.

³Kiyotaki and Wright assume random pairings of agents, and define an equilibrium of the system as a state in which the fraction of agents holding each type of good is constant over time. This requires an infinite population and an (implicit) application of a law of large numbers that allows replacing payoffs by their expected values.

few bytes long, and starting from a random distribution of genotypes, the simulations invariably converge to the appropriate equilibria.⁴

Each agent in our economy has an *exchange strategy* specifying which trades are acceptable from a given pre-trade inventory, and a *production/consumption strategy* specifying what the agent consumes and produces from a given post-trade inventory. Agents meet randomly in each period with their pre-trade inventories, they deploy their exchange strategies to identify a mutually acceptable trade, and they then produce and consume, given their post-trade inventories, according to their production/consumption strategies.⁵

We make several assumptions that allow us to represent this economy as a Markov process. First, we assume agents cannot make binding contracts across time periods. Second, we assume the economy is ‘very large,’ so an agent’s history of trades gives no information concerning the agent’s future trading opportunities.⁶ Third, we assume agents do not recognize one another as individuals in distinct trading periods, so there is no scope for the use of strategies that take into account the prior history of trade of individual agents. This assumption reflects our intention of modeling anonymous exchange rather than repeated interactions. These assumptions together imply that an agent’s strategy depends only on the agent’s current state (and perhaps time).

Finally, we assume each agent has a constant probability ρ of living to the next period, and an agent who dies is immediately replaced by an identical agent with an empty inventory.⁷ Since agents do not age (i.e., their

⁴The exception to this rule is in the case of *fiat* money (defined below), in which case we expect on general grounds a rather small basin of attraction for the monetary equilibrium.

⁵We could readily extend this model by assuming each agent has a finite set of search strategies and the probability of two agent’s meeting in is a function of the search strategies of all the agents in the economy. This would allow us to define ‘marketplaces,’ so that agents with particular production goods and consumption needs can locate one another with higher frequency and at lower cost than by random pairing, as well as ‘buyers’ and ‘sellers,’ so that something might emerge approaching the empirically observable fact that agents alternate between strategies of buying and selling (?).

⁶Since we are interested in the ergodic properties of a monetary economy, we do not model learning. Rather, agents are assumed to know the structure of the economy and can infer from this the probability distribution over states of the economy. Moreover, the ergodic assumption implies that the personal histories of an agent does not affect the agent’s assessment of these probabilities. The simulations of course make no such assumption.

⁷To ensure that there are a finite number of agents despite the fact that agents die and

probability of death is independent of the number of periods they have been alive) and since at birth each agent assumes the economy is in its long run ergodic state, agents' strategies can be assumed time-independent. Finally, we assume that aside from their personal characteristics and holdings, agents know only the ergodic distribution over states in the economy, on the basis of which they can calculate the payoffs to alternative trading strategies. These assumptions allow us to conclude that if we define the 'current state' of the economy as the set of pre-trade inventories of agents, the next state of the economy depends only on its current state and the agents' choice of strategies. Moreover, agent strategies are time-stationary.

We prove that for each choice of a trading strategy by each agent, the resulting stochastic process has an ergodic distribution of states. We assume agents choose strategies with maximal lifetime payoffs subject to this ergodic distribution of states. Since the equilibrium distribution of states is itself a function of the strategies chosen by the agents, the natural concept of equilibrium in the choices of agents is that of Nash equilibrium.⁸

Among the weaknesses of standard general equilibrium theory is the absence of any economically meaningful dynamic leading to the stability of equilibrium (?). Evolutionary models, of which ours is an example, do not have this weakness. Indeed, we may posit a dynamic much as in standard evolutionary theory, in which strategies that pay off for the agents who use them tend to increase in frequency over time at the expense of strategies with poor payoff records. In Section 5, we construct such a dynamic economy, in which artificial life creatures are endowed with a genetic structure encoding

are born, we assume a newly born agent inherits the strategy of its predecessor. This is a harmless assumption, since if agents use Nash strategies, a new agent can do no better than adopt the strategy of its predecessor. Moreover we could replace the assumption that agents die with the assumption that with a positive probability, however small, they lose their inventory. We prefer the 'death' assumption, since it extends naturally to the dynamic simulation, in which agent 'inherit' strategies from their predecessors, and successful strategies 'evolve' by having more representatives in succeeding generations.

⁸Evolutionary stability rather than Nash might be considered more appropriate for a biologically-inspired model. However since all goods have positive storage costs in our model, a mutant could not profitably change its choice of optimal inventory in anticipation of meeting a mutant trading partner, since if the probability of such a meeting is sufficiently small, the costs must outweigh the expected gains. Evolutionary stability would rule out some unreasonable trading strategies, but these would be ruled out by dropping Pareto-dominated Nash equilibria as well. Thus there would be little gain in requiring evolutionary stability. A more appropriate refinement of Nash equilibrium would be local or global stability under a replicator dynamic. We explore such a dynamic in the simulations below.

particular trading strategies.⁹

Previous artificial life simulations of a Kiyotaki-Wright model (?), using the classifier-based genetic algorithms of John Holland (holland:adaptation, holland:brittleness), have shown themselves to have a very narrow window of convergence, and to be extremely sensitive to initial conditions and specifications of the parameters of the model.¹⁰ The simulations reported here, by contrast, are based on genetic algorithms that implement a Darwinian evolutionary dynamic and are robust under a wide variety of specifications. As an indication of robustness, whereas simulations based on classifier systems have been restricted to a three-good economy, our algorithms allow us easily to simulate a five-good economy, and the quality of the simulations indicates the feasibility of modelling larger numbers of goods as well.¹¹

2 General Equilibrium with Bilateral Trade

Let G and A be finite sets with at least two elements. G represent the set of goods in the economy, and A represents the set of agents. Each agent $a \in A$ consumes a nonempty set of goods $C_a \subset G$ and produces a nonempty set of goods $P_a \subset G$, where $C_a \cap P_a = \emptyset$.¹² Each agent $a \in A$ also has an increasing, non-negative and bounded utility function $u_a : Z_+[C_a] \rightarrow \mathbf{R}$, where \mathbf{R} is the real numbers and $Z_+[S]$ is the set of non-negative integral multiples of elements of the set S .¹³ Each good $g \in P_a$ also has a production cost $\kappa_{ag} > 0$,

⁹The idea of treating cultural units as replicators in an evolutionary dynamic was first suggested by (?). For evolutionary models of cultural development, see (?), (?), and (?). The importance of evolutionary dynamics has also been stressed by (?).

¹⁰I replicated the Marimon, McGrattan, and Sargent's simulations using a version of Riolo's riolo86 CFS-C classifier program adapted for Turbo C running under MSDOS, and found the model unacceptably erratic and sensitive to the choice of parameters.

¹¹For a general introduction to genetic algorithms, see (?). The algorithms used in this paper were written by the author in Borland Pascal for the MSDOS operating system, and are available from my home page <http://www-unix.oit.umass.edu/~gintis>.

¹²Note that production does not requires only labor inputs. An interesting generalization would be to allow for production inputs. In this case we would have to introduce a price system, and a necessary condition for a trading equilibrium would be that newly born agents have enough goods in inventory to produce something for trade.

¹³In a model of face-to-face trading is more reasonable to assume that the agent's utility function and production possibilities are a function of the agent's current *state*, and that the current state is a function of past consumption and production decisions. Such an assumption also allows us to represent diminishing marginal utility, complementarity in

and we assume for each $c_a \in C_a$, $u_a(c_a) > \min_{g \in P_a} \kappa_{ag}$ (i.e., the agent has at least one production good whose cost of production is less than the utility of consuming c_a). We extend κ_a linearly to a function $\kappa_a : Z_+[P_a] \rightarrow \mathbf{R}$ on the set of production goods such that for $p = \sum_{g \in P_a} n_{ag}g$,

$$\kappa_a(p) = \sum_{g \in P_a} n_{ag} \kappa_{ag}.$$

In addition, each agent $a \in A$ incurs a storage cost $s_{ag} > 0$ in each trading period for each unit of a good g held in inventory. Since all goods are indivisible, we can identify inventories with subsets of $Z_+[G]$. We extend s_{ag} linearly to a function $s_a : Z_+[I_a] \rightarrow \mathbf{R}$ of the set I_a of inventories, so that for $i_a = \sum_{g \in G} n_{ag}g$,

$$s_a(i_a) = \sum_{g \in G} n_{ag} s_{ag},$$

The utility accruing to agent $a \in A$ at date t is then

$$v_{at} = u_a(c_{at}) - \kappa_a(p_{at}) - s_a(i_{at}), \quad (1)$$

where $c_{at} \in Z_+[C_a]$ is the set of goods consumed at date t , $p_{at} \in Z_+[P_a]$ is the set of goods produced at date t , and i_{at} is the inventory the agent will take into the next trading period.¹⁴ The total return to agent a is then

$$v_a = \sum_{t=0}^T v_{at}, \quad (2)$$

where T is a random variable representing the period in which the agent dies.

We call the vector $[C_a, P_a, \kappa_a, s_a, u_a]$, where the entries are defined above, the *type* of the agent, and we assume throughout that if the economy has one agent of a given type, it has at least two agents of that type. We also assume

consumption, and the use of factor inputs in production. Our analysis extends directly to this case, provided there are a finite number of agent states (we then must include the state of each agent, in addition to the state of the agent's inventory, in the specification of the state of the economy). For notational simplicity, we assume there is only one state for each agent.

¹⁴At the cost of some additional notation, we could allow agents to be self-sufficient by including a production process for each 'necessity' an agent consumes. In this case the utility of consumption of a necessity should be replaced by its cost of own-production. The emergence of trade in such a model entails specialization in production in addition to higher agent welfare.

each good is produced by at least one agent and each good is consumed by at least one agent unless otherwise stated.

For agent $a \in A$, we call an inventory i_a *admissible* if its total storage costs $s_a(i_a)$ are less than

$$\sup\{u_a(\tilde{c}) \mid \tilde{c} \in Z_+[C_a]\}/\rho,$$

which is finite (here ρ is the probability that an agent survives to the next period). Since storage costs are strictly positive, the set I_a of admissible inventories for agent a is finite. An optimizing agent never moves from an admissible to an inadmissible inventory, since the agent would certainly do better simply to discard the whole inventory and start over. We thus lose no generality by assuming that there is some n_{max} such that no agent's inventory can contain more than n_{max} items.

For simplicity, we assume that one encounter between randomly paired agents occurs in each trading period. A pure *trade strategy* $\sigma_a = [\succeq_a, c_a, p_a]$ for agent $a \in A$ is a total ordering \succeq_a on I_a , a function $c_a : I_a \rightarrow Z_+[C_a]$ and a function $p_a : I_a \rightarrow Z_+[P_a]$. Here $i_a \succeq_a j_a$ means agent a weakly prefers trading inventory $i_a \in I_a$ for inventory $j_a \in I_a$, rather than holding i_a ; $c_a(i_a) \subseteq i_a$ is the vector of goods consumed from post-trade inventory i_a ; and $p_a(i_a)$ is the vector of goods produced when the post-trade inventory is i_a .¹⁵ The set Σ_a of pure trade strategies is clearly finite and non-empty. We write $\Sigma = \prod_{a \in A} \Sigma_a$, which is the set of pure trade strategies for the whole economy.

Let $a, b \in A$ be distinct agents with trade strategies $\sigma_a = [\succeq_a, c_a, p_a]$ and $\sigma_b = [\succeq_b, c_b, p_b]$ holding admissible trade inventories i_a and i_b , respectively. The agents' *trading set* $G(i_a, i_b; \succeq_a, \succeq_b) \subseteq I_a \times I_b$ is the set of inventories obtained through trades that weakly dominate (i_a, i_b) . That is $(j_a, j_b) \in G(i_a, i_b; \succeq_a, \succeq_b)$ iff $j_a \cup j_b = i_a \cup i_b$, $j_a \succeq_a i_a$, and $j_b \succeq_b i_b$. $G(i_a, i_b; \succeq_a, \succeq_b)$ is non-empty, since it contains (i_a, i_b) when the agents do not gain from trade. If $G(i_a, i_b; \succeq_a, \succeq_b)$ contains more than one element, we assume a trade is chosen randomly from the Pareto efficient elements in $G(i_a, i_b; \succeq_a, \succeq_b)$.¹⁶

¹⁵We assume that agents can freely discard goods in their inventory by 'consuming' them with a utility payoff of zero.

¹⁶It would be more reasonable to treat this situation as a bargaining game, as developed for instance in (?). This would also allow considerations of excess demand and supply to affect the terms of trade, since an agent's fallback position in case no agreement is reached is stronger, the relative scarcity of the goods the agent sells, and the relative abundance of

Let $I = \prod_{a \in A} I_a$ be the set of admissible inventories for the economy, and for $i \in I$, we write $i[a]$ for the inventory of agent a in i . Let $\sigma = \{\sigma_a \mid a \in A\}$ be a choice of a pure trade strategy for each agent. We call σ a *pure strategy* for the economy. Given strategy σ , the economy becomes a Markov chain \mathbf{M}^σ with transition probability matrix $P^\sigma = \{P_{ij} \mid i, j \in I\}$, in which the current state is the vector of post-consumption, pre-production agent inventories. An agent $a \in A$ who dies is replaced by an identical agent who inherits a 's label and representation in I , with an empty inventory but the possibility of producing in anticipation of the next trading period. The Markov property is satisfied for this system, since the probability distribution over next-period states depends on the array of post-consumption inventories in the current period alone, given σ .

The Markov chain M^σ is aperiodic, and the set of states accessible from the ‘zero state,’ in which all agents have empty inventories, is an irreducible recurrent class. To see this, note that given σ , since each agent has a nonzero probability of dying, every state of \mathbf{M}^σ has a nonzero probability of transition to the zero state. It follows that the zero state is recurrent, and the class of states attainable from the zero state is an irreducible recurrent class. This class, moreover, is aperiodic and contains all recurrent states, since every state in the class has the same period, and every state has a nonzero probability of transition to the zero state.

To define a concept of equilibrium for the economy, let Ω be the sample path space generated by I ; i.e., $\omega \in \Omega$ is a sequence

$$\omega = (\omega_1, \omega_2, \dots, \omega_t, \dots) \in \prod_{t=0}^{\infty} I,$$

where ω_t is the post-trade inventory at date t . Let us write ω_{at} for the component of ω_t corresponding to agent a . From (1), the flow of utility accruing to an agent $a \in A$ in period t from path $\omega \in \Omega$ is then

$$v_{at}(\omega; \sigma) = u_a(c_a(\omega_{at})) - \kappa_a(p_a(\omega_{at})) - s_a(\omega_{a,t+1}), \quad t \geq 1 \quad (3)$$

and for $t \geq 1$,

$$\omega_{a,t+1} = \begin{cases} \emptyset & \text{with probability } \rho \\ \omega_{at} - c_a(\omega_{at}) + p_a(\omega_{at}) & \text{with probability } 1 - \rho \end{cases}$$

the goods the agent buys. We ignore this alternative here, since we are not here concerned with the process of relative price determination.

By (2), the payoff to player $a \in A$ with trade strategy σ_a and trading history $\omega \in \Omega$ can be written as

$$v_a(\omega, \sigma) = \sum_{t=0}^T v_{at}(\omega; \sigma_a). \quad (4)$$

where T is a random variable representing the agent's death. The payoff to agent a is then given by the expected value of (4) with respect to the appropriate probability measure over the paths $\omega \in \Omega$.¹⁷

We now define the payoff to agent $a \in A$, given that the initial state of the economy is i_0 by

$$v_a(\sigma; i_0) = \mathbf{E}\{v_a(\omega; \sigma) \mid \omega_0 = i_0\} \quad (5)$$

where the expectation is with respect to $\pi[\sigma, i_0]$.

The ergodic theory for finite Markov chains (?):91 ensures that an invariant measure (a 'long run equilibrium') for this economy exists. According to the ergodic theorem, given σ there is a unique probability distribution π^σ over I such that

- (a) If the states in I are distributed as π^σ in period t , they remain distributed as π^σ in period $t + 1$; i.e.,

$$\pi_j^\sigma = \sum_{i \in I} \pi_i^\sigma P_{ij}^\sigma \quad j \in I;$$

- (b) For all states $i, j \in I$, starting in state i at $t = 0$, the probability of being in state j at time t approaches π_j^σ as $t \rightarrow \infty$; i.e.,

$$\pi_j^\sigma = \lim_{t \rightarrow \infty} P_{ij}^t \quad i, j \in I;$$

¹⁷To define this measure, for $n > 0$ let $\Omega^n = I \times \dots \times I$, where the product is taken n times. Then for each initial state i_0 , we define a probability distribution $\pi_n[\sigma, i_0] : \Omega^n \rightarrow \mathbf{R}$ such that the probability of a path $\omega = (\omega_1, \dots, \omega_n) \in \Omega^n$ is given by

$$\pi_n[\sigma, i_0](\omega) = P_{\omega_0 \omega_1}^\sigma \cdots P_{\omega_{n-1} \omega_n}^\sigma,$$

which is of course the probability that the Markov chain M^σ follows path ω on the first n transitions starting from initial inventory i_0 . By the Kolmogorov Consistency Theorem (?):189, the sequence of distributions $\{\pi_n[\sigma, i_0] \mid n = 1, 2, \dots\}$ extends uniquely to a probability distribution $\pi[\sigma, i_0]$ on Ω .

Thus given σ and an initially empty inventory for each agent, after an initial ‘startup period,’ the Markov economy settles down into its ergodic condition in which inventories are distributed according to π^σ . We thus take the payoff to each agent a as the expected return to σ_a when a starts with an empty inventory and the other agents’ inventories are distributed according to the ergodic probability distribution π^σ . Specifically, let π_a^σ be the conditional distribution on Ω given by $\pi_a^\sigma(i) = \pi^\sigma(i \mid i[a] = \emptyset)$. Then we define

$$v_a(\sigma) = \mathbf{E}\{v_a(\omega; \sigma)\} \tag{6}$$

where the expectation is with respect to the probability distribution π_a^σ , which is the expected payoff to agents following strategy σ when the information of each agent is just the transition probability matrix P^σ .

We assume the initial inventory i_0 is a draw from the ergodic distribution for σ . We could assume i_0 is the zero state (an empty inventory for all agents), but then it is not reasonable that agents choose a stationary trade strategy, and we cannot justify the assumption that newly born agents take the same trade strategies as the agents they replace. Rather, it is reasonable to assume that each agent considers the inventories of other agents in the economy to be in some sense ‘in long-run equilibrium’ with respect to trade strategy σ .

Having defined the payoffs to each agent for each trade strategy σ , our Markov economy is a finite game. For each $a \in A$, let Σ_a^* be the set of mixed strategies for a , and let Σ^* be the corresponding set of mixed strategies for the Markov economy as a whole. If agents adopt mixed strategy $\sigma^* \in \Sigma^*$, the probability $P_{ij}^{\sigma^*}$ of passing from state i to state j given σ^* is simply the appropriate convex combination of the P^σ ’s for the pure strategies $\sigma \in \Sigma$ underlying σ^* , and payoffs to the agents are the appropriately weighted sum of the payoffs to the underlying pure strategies. By Nash’s theorem there is at least one Nash equilibrium in mixed strategies.

But of course this result is trivial, since the strategy of each agent not producing and not trading is clearly a Nash equilibrium. We call this the *no-trade equilibrium*. However we have the more useful

Theorem 1 *Suppose each agent $a \in A$ is constrained to a nonempty pure strategy set $\Upsilon_a \subseteq \Sigma_a$. Then there exists a Nash equilibrium of the Markov economy in mixed strategies based on the pure strategies in $\{\Upsilon_a \mid a \in A\}$.*

This theorem is useful because in some cases we can find interesting Nash equilibria by first constraining the pure strategies of the agents, and then

showing that when we relax the constraint the equilibrium remains Nash. For instance we have

Theorem 2 *Suppose each agent produces exactly one good, and for any two goods g and h there is a producer of g that consumes h . If the utility of consumption is sufficiently great for each good, there is a Nash equilibrium in which all goods are traded with positive probability.*

Proof: We first constrain the pure strategies of the agents as follows. We assume each agent must hold a unit of some good in its inventory in each period. Also, an agent must offer to trade its production good p_a unit for unit against any consumption good $c_a \in C_a$. By Theorem 1 there exists a Nash equilibrium σ for the resulting ‘constrained’ Markov economy.

We show that if the utility of consumption is sufficiently great for each good, there is a positive probability of trade in each period in the constrained Nash equilibrium. Suppose for each agent, the utility of a unit of each consumption good for that agent exceeds the cost of production for that agent. Suppose there is some good, say g , that is not traded. Let $a \in A$ be a producer of g , so a holds g in inventory in each period. Suppose a consumes h , and let $b \in A$ be an agent who produces h and consumes g . If h is traded with zero probability then b surely holds h in inventory, and since the probability of a encountering b is nonzero, the probability of a trade of g for h is nonzero, which is a contradiction. Thus h is traded with positive probability. But then after a successful trade, b must produce a unit of h , and with positive probability b will encounter a in the next period, in which case a trade will occur. The probability of a trade of g for h is therefore strictly positive.

Now let us drop the restrictions on the strategies of the agents in the economy. To show that the strategy σ remains Nash when consumption utilities are sufficiently large, we choose an arbitrary agent $a \in A$, we assume all other agents follow $\sigma_{-a} = \{\sigma_b \mid b \in A, b \neq a\}$, and we show that σ_a remains a best response to σ_{-a} . The only way σ_a could fail to be a best response is if either σ_a is dominated by the no-trade strategy, or if it remains optimal for a to produce its production good p_a , but there is no $c_a \in C_a$ that it is optimal to accept in trade for p_a . The latter possibility is ruled out by the one-stage deviation principle (?):109, assuming the utility of consumption is sufficiently large (i.e., assuming $u_a(c_a) > \kappa_{ap_a}$).

So suppose it is optimal for agent $a \in A$, who produces p_a and consumes c_a , to follow the no-trade strategy, which has zero payoff. Then a ’s strategy τ_a of producing p_a and holding it in inventory until a encounters an agent

who holds c_a in inventory and is willing to accept p_a in trade must have a negative payoff. Let $b \in A$ produce $c_a \in C_a$ and consume p_a . Since there is another agent of a 's type in the economy, say a' , there is a positive probability that b will meet a' in each period, in which case b will enter into an exchange with a' , after which b will produce c_a and hold it in inventory. Thus a has a positive probability, say β , of meeting b carrying c_a in each period. The return to τ_a must then satisfy the recurrence relation

$$\begin{aligned} v_{\tau_a} &\geq \beta(u_a(c_a) - \kappa_{ap_a}) - s_a(p_a) + (1 - \rho)v_{\tau_a} \\ &= \beta(u_a(c_a) - \kappa_{ap_a}) - s_a(p_a)]/\rho, \end{aligned}$$

which is positive for sufficiently large $u_a(c_a)$. This is a contradiction, which proves Theorem 2.

3 The Existence of a Medium of Exchange

Definition 1 *We say a good $m \in G$ is a universal medium of exchange if there is a positive probability of trade in each period and m is accepted in trade by all agents against all other goods $g \neq m$. We say a good $m \in G$ has low storage costs if for each $a \in A$, $s_{am} \leq s_{ap_a}$, where $p_a \in P_a$ (i.e., the storage cost of m is not greater than the storage cost of a 's production goods).*

Theorem 3 *Suppose each agent $a \in A$ produces exactly one good p_a , and let $m \in G$ have low storage costs. Suppose that for each $g \in G$ there is an agent who produces m and consumes g . Then if the utility of consumption is sufficiently high, there is a Nash equilibrium of the economy in which m is a universal medium of exchange.*

Proof: Suppose $m \in G$ satisfies the conditions of the theorem. We add three constraints to agent's pure trade strategies. First, we constrain each agent to hold a unit of some good in its inventory in each period. Second, for each agent $a \in A$, we admit only those trade preferences \succeq_a for which a 's post-trade inventory is no smaller than a 's pre-trade inventory. This implies, of course, that if trade occurs relative prices for all goods are unity. Third, we require all agents to choose strategies in which one unit of m is accepted in exchange for each good $g \neq m$ that they currently hold in their inventories. By Theorem 1, there exists a Nash equilibrium for this economy. To show that m is a universal medium of exchange in this constrained economy, we

must show that there is a positive probability of trade. We first note that every good will be offered for sale with positive probability, since for each good $g \in G$ there is an agent a who produces g , and with positive probability a is replacing an agent of the same type who died in the previous period, and hence a must be holding g in inventory. But if the probability of trade is zero, then all agents must be holding m in inventory with probability one, and since there are at least two goods, this is impossible. Thus there is a positive probability of trade in each period. Indeed, each good will trade with positive probability, since for any good g with positive probability there is an agent holding g in inventory, and that agent will encounter an agent willing to accept g in trade for m (namely, any agent producing m and consuming g).

We must show that each agent's strategy will remain unchanged when these constraints are relaxed for this agent but remain for all other agents. We first show that for each agent $a \in A$, σ_a dominates the no-trade strategy if the utility of consumption is sufficiently large. We begin by assuming that a produces $p_a \neq m$. Consider the following constrained strategy τ_a : with an empty inventory, produce one unit of p_a ; exchange p_a for either m or a consumption good $c_a \in C_a$. Since there is some agent who produces m and consumes p_a , a will encounter an agent willing to trade p_a for m with positive probability. In the next period a will encounter an agent selling some $c_a \in C_a$ with positive probability, since each good is offered for sale with positive probability (we use the assumption that there are at least two agents of each type here). Thus with positive probability the agent will trade p_a for some $c_a \in C_a$ in two periods, and if $u_a(c_a)$ is sufficiently large, strategy τ_a strictly dominates the no-trade strategy, and since σ_a weakly dominates τ_a , σ_a strongly dominates the no-trade strategy, and hence must involve trade with positive probability.

Now suppose a produces m . Since a meets an agent holding some $c_a \in C_a$ with positive probability, a will trade with positive probability, and hence if the utility of consumption is sufficiently high, trading dominates the no-trade strategy for an m -producer as well.

Now for any $a \in A$, since a does not use the no-trade strategy, a 's inventory will never be empty (we use the one-shot deviation principle here). For the same reason a must continue to offer to trade at least one-for-one. Of course a could offer many-to-one trades, but these are weakly dominated by one-to-one trades, since we assume free disposal. Now suppose a holds good g in inventory and encounters an agent willing to accept g in exchange for

m. We must show that *a* remains willing to accept *m* in exchange for *g*. If *a* does not wish to consume *m*, since $s_{ag} \geq s_{am}$ and every agent who accepts *g* also accepts *m*, the payoff to holding *m* in inventory must be at least as great as that of holding *g* in inventory. Hence *a* will accept *m* in trade in this case as well. Note that this argument holds even in case *a* is an *m*-producer. This completes the proof of Theorem 3.

4 Simulating an Economy: Memes, Genetic Algorithms, and Replicator Dynamics

Genetic algorithms and other artificial life approaches to complex dynamics have been largely inspired by contemporary population biology, which is in turn indebted not only to Darwinian notions of evolution but to the genetic details specific to DNA-based life forms. Sexual reproduction, diploidy, recombination, crossover, dominance, meiotic drive, and related notions, however, were not part of the conceptual framework of evolutionary theory until relatively recently. Moreover, it is likely that such mechanisms are not necessary elements in the emergence of life, but rather particular to the mechanics of DNA replication and the chemistry of terrestrial life. There is no particular reason to mimic such phenomena in artificial life. We therefore abandon them as inappropriate to modeling an economy.

On the other hand, Darwin's basic notions of natural selection, random variation (mutation), and adaptation are central to the evolutionary model, and can be made to apply to many phenomena in which the mechanics of replication have little in common with DNA-based mechanisms.¹⁸ In particular, we can think of 'culture' as a heterogeneous collection of behavioral mechanisms (following ?) we will call them 'memes'), which human agents draw upon in their social interactions. The frequency of a particular meme is then the fraction of the population that deploys it, as opposed to one of its 'alleles' (alternative mechanisms capable of being deployed under the same circumstances). Many authors have proposed evolutionary cultural transmission mechanisms, following the seminal work of (?), (?), and (?).

¹⁸In fact something akin to 'crossover' is also doubtless important in modeling cultural evolution, in the form of the ability of agents to assume part of the repertoire of behavior of another agent rather than the whole. In our simulation there is essentially only one genetic locus, so we do not model this phenomenon.

I will show here that a transmission mechanism that is quite plausible in the context of economic transactions leads to the standard replicator dynamic of evolutionary theory. I will then use this dynamic in the rest of the paper for simulating the emergence of money in an artificial life community.¹⁹ The mechanism requires that in each period, each agent have imperfect knowledge of how well his current strategy is performing, as well as that of one other agent in the economy with whom he has randomly come into contact (e.g., a trading partner). With a certain probability an agent abandons his current strategy if the other agent's strategy appears to be doing better.

Suppose the economy has many agents, each following a pure strategy s_i for $i = 1, \dots, n$. The play is repeated in periods $t = 1, 2, \dots$. Let x_i^t be the fraction of players playing s_i in period t , and suppose the payoff to s_i is π_i^t . Without loss of generality, for a given period t we can assume $\pi_1^t \leq \pi_2^t \leq \dots \leq \pi_n^t$.

Suppose in every period, with probability $\alpha > 0$, an agent using s_i compares his strategy with a randomly chosen other agent, and changes to the other's strategy if he perceives that its payoff is higher. However information concerning the difference in the expected payoffs of the two strategies is imperfect, so the larger the difference in the payoffs, the more likely the agent is to perceive it, and change. Specifically, we assume the probability p_{ij}^t of shifting from s_i to s_j is given by

$$p_{ij}^t = \begin{cases} \beta(\pi_i^t - \pi_j^t) & \text{for } \pi_i^t > \pi_j^t; \\ 0 & \text{for } \pi_i^t \leq \pi_j^t; \end{cases}$$

The expected fraction $\mathbf{E}[x_i^{t+1}]$ of the population using s_i in period $t + 1$ is then given by

$$\begin{aligned} \mathbf{E}[x_i^{t+1}] &= x_i^t - \alpha x_i^t \sum_{j=i+1}^n x_j^t \beta(\pi_j^t - \pi_i^t) + \sum_{j=1}^i \alpha x_j^t x_i^t \beta(\pi_i^t - \pi_j^t) \\ &= x_i^t + \alpha x_i^t \sum_{j=1}^n x_j^t \beta(\pi_i^t - \pi_j^t) \\ &= x_i^t + \alpha x_i^t \sum_{j=1}^n \beta(\pi_i^t - \pi_j^t) x_j^t \\ &= x_i^t + \alpha x_i^t \beta(\pi_i^t - \bar{\pi}^t), \end{aligned}$$

¹⁹In fact the simulations below combine a replicator dynamic with Darwin's 'random variation' mechanism, in that copying of strategies is subject to random mutation.

where $\bar{\pi}^t = \pi_1^t x_1^t + \dots + \pi_n^t x_n^t$ is the average return for the whole population. If the population is large, we can use the law of large numbers to replace $\mathbf{E} [x_i^{t+1}]$ by x_i^{t+1} , and write $\dot{x}_i^t = x_i^{t+1} - x_i^t$, getting

$$\dot{x}_i^t = \alpha \beta x_i^t (\pi_i^t - \bar{\pi}^t), \quad \text{for } i = 1, \dots, n.$$

This is the *replicator dynamic* of evolutionary theory (?).

5 Constructing an Artificial Life Economy

We construct an economy with five goods g_1, \dots, g_5 , ordered from highest to lowest in storage costs. We assume that (unless otherwise stated) for each pair of goods, there is a type of agent in the economy producing the first good and consuming the second. There are thus twenty agent types, since we assume that for each good g , the utility of consumption u_g , the cost of production κ_g , and the storage cost s_g are the same for all agents who consume, produce, and store g , respectively. For simplicity we also assume that an agent's maximum inventory is one unit. This ensures that all trades are one-for-one, thus speeding up the simulations. Since we are not concerned here with the formation of relative prices, this appears to be a reasonable restriction.

An individual agent's *genotype* uniquely defines its trading strategy. Genotypes are assigned randomly to the initial population, and inherited (with possible mutation) by agents born in later periods. We represent the genotype by twenty five bits of computer memory, which we map into a 5×5 matrix as follows,

$$\underbrace{[abcdefghijklmnopqrstuvwxy]}_{\text{agent genome}} \longrightarrow \begin{matrix} & g_1 & g_2 & g_3 & g_4 & g_5 \\ g_1 & \left(\begin{array}{ccccc} a & b & c & d & e \\ f & g & h & i & j \\ k & l & m & n & o \\ p & q & r & s & t \\ u & v & w & x & y \end{array} \right) \end{matrix}$$

where each of the letters is 1 or 0. Thus, for instance, if b is 1, then the agent with this genome will trade one unit of g_1 for one unit of g_2 , should the agent have g_1 in its inventory and encounter a trading partner holding g_2 and willing to trade for g_1 . Similarly, if n is 0, the agent will not trade a unit

of g_3 for a unit of g_4 even if the agent has g_3 in its inventory and encounters a willing trading partner.

In our general model agents have production and consumption strategies in addition to trading strategies. We fix these strategies, however, by constraining each artificial life agent to consume when it has acquired its consumption good, and to produce a unit of its production good when its inventory is empty. These assumptions are not substantive in this context, and merely speed up the rate of convergence of the system to equilibrium. If we were to allow agents to follow a no-trade strategy, in the long run either the no-trade agents would disappear or the whole economy would settle into a no-trade equilibrium (which of the two occurs depends on the number of agents who trade in the initial draw of the population). If we allow agents to hold their consumption good without consuming it, agents who do so are eliminated in very few rounds. Thus no interesting behavior is bypassed through these restrictions.

With five goods there are $2^{25} = 33,554,432$ distinct trading strategies. We can reduce this number to $2^{16} = 65,536$ by eliminating dominated strategies (for instance strategies that involve not accepting a consumption good in trade, or trading a good for itself). Equating operationally-equivalent strategies (for instance strategies that differ on what to accept in trade when an agent sells her consumption good, which can never occur) reduces this to $2^{12} = 4,096$ distinct strategies. Since there are twenty types of agents, and assuming one player of each type, there are $4096^{20} \approx 10^{72}$ pure strategies for the economy, an astronomically large figure. However using genetic algorithms, this strategy space can be efficiently searched in a short time using a total of 1000 to 2000 agents (50 to 100 agents of each of the twenty types). Repeated runs of the model with different random draws of the population on each run show that the model converges to roughly the same equilibrium after about 6,000 generations (which is about 60,000 trading sessions in our model).

A simulation begins by creating a population of agents with randomly generated genomes, empty inventories and zero wealth. In each round of trade, each agent chooses a good from inventory to sell, and trading partners are assigned randomly. If they agree to trade, each agent's wealth is augmented by u_c if the agent acquires its consumption good, and decremented by k_p if it must produce its production good. Whether or not trade takes place, each agent's wealth is debited an amount equal to the storage costs on the goods it holds in inventory.

Every ten rounds, which we call a ‘generation,’ some agents die and others are born. The least fit agent of each type dies, and is replaced by the offspring of one of the top-performing agents of this type, and the offspring is endowed with the average wealth for that type of agent. In addition, elderly agents (defined as those who have lived for more than ten generations) die at a given rate and are replaced by offspring of the more successful agents. Reproduction takes the form of asexual replication with mutation (one gene locus is randomly chosen to mutate for each new agent).²⁰

Since mutation is constantly taking place, it is clear that this model can never sustain optimality, since optimal strategies have a positive probability of mutating into suboptimal strategies in each generation. I used two criteria to specify when agents of a particular type have attained an equilibrium trading strategy: (a) At least 80% of agents of this type, and the wealthiest 35% of agents of this type, accept the same goods in trade for their production good.²¹

I will report on a typical simulation, in which 100 agents of each of the twenty types were created, and the storage costs for the five goods (g_1, g_2, g_3, g_4, g_5) were set equal to (40,30,20,10,2). According to Theorem 3, we expect g_5 to serve as money for all agents (except of course g_5 -producers and -consumers), as long as the utility of consumption is sufficiently high. Moreover, it is reasonable to expect that g_4 serve as money for g_1, g_2 , and g_3 -producers who do not consume g_4 , that g_3 serve as money for g_1 and g_2 -producers who do not consume g_3 , and that g_2 serve as money for g_1 -producers.²²

²⁰Genetic algorithms generally allow for more flexible searching of the decision space by including sexual reproduction and crossover. These mechanisms were not required in our model, since the performance of the simpler parthenogenetic algorithms performed quite adequately.

²¹This criterion assumes that all agent types use pure strategies. If a mixed strategy is optimal, the criterion may not pick this up. In fact, in almost all simulations, almost all types did stabilize according to this criterion. There is a way to detect mixed strategy equilibria, using that fact that under plausible conditions an evolutionary model with a replicator dynamic (agents with higher fitness having higher reproduction rates) will exhibit oscillatory behavior around the equilibrium frequencies rather than approaching equilibrium asymptotically. I did not test formally for such behavior, but I did not observe it in any simulation.

²²Of course Theorem 3 asserts only the *existence* of an equilibrium of this type, not that there is an evolutionary dynamic tending to this equilibrium. Our simulations indicate that there is such a dynamic.

The simulation ran for 553 generations, or 5530 rounds of trading, before all agent types stabilized, in the sense that the frequencies of the various strategies remained constant over time except for the random perturbations induced by mutation. The resulting pattern of good-acceptance is shown in Table 1. The results compare well with our expectations. Good g_5 is a universal medium of exchange, g_2 and g_4 are accepted in trade by all types whose storage costs exceed that of g_2 and g_4 , respectively. Good g_3 , however, is not accepted by agents who produce g_2 and consume g_1 , nor is it accepted by agents who produce g_1 and consume g_2 or g_4 .²³

Agent Consumes	Agent Produces				
	g_1	g_2	g_3	g_4	g_5
g_1		g_1, g_4, g_5	g_1, g_4, g_5	g_1, g_5	g_1
g_2	g_2, g_4, g_5		g_2, g_4, g_5	g_2, g_5	g_2
g_3	g_2, g_3, g_4, g_5	g_3, g_4, g_5		g_3, g_5	g_3
g_4	g_2, g_4, g_5	g_3, g_4, g_5	g_4, g_5		g_4
g_5	g_2, g_3, g_4, g_5	g_3, g_4, g_5	g_4, g_5	g_5	

Table 1: **The Emergence of a Commodity Money Equilibrium.** Shows the goods agents accept in trade for their production good, for all agent types, with storage costs (40,30,20,10,2). Note that in equilibrium, each type of agent tends to accept goods with lower storage costs than its production good, and g_5 is universally accepted as a medium of exchange.

Table 2 gives a more disaggregated view of one type of agent, the producer of g_2 and consumer of g_4 , from the random distribution of strategies in generation 0, through intermediate stages in generations 50, 100, and 200, to its stabilization at generation 314. Note that Table 2 is still somewhat aggregated, since it distinguishes trading strategies only according to what goods are accepted in trade for money (g_5) and the agents' production good (g_2). Table 2 shows that strategies that do not accept g_5 as money are virtually eliminated by generation 100, and strategies that are willing to trade the production good g_2 for a higher-storage cost good (g_1) do show some oscillation before being eliminated.

²³Other simulations conformed to this pattern, including at times the tendency of g_3 not to be accepted by agents who would appear to benefit by doing so.

Trading Round	Buys Non- c with Money	Buys Higher Cost Good with p	# of Distinct Strategies Used
0	81	54	30
500	47	48	22
1000	7	2	8
2000	0	15	4
3140	0	0	3

Table 2: **Evolution Towards an Equilibrium Trading Strategy:** Shows the progress of the 100 agents producing g_2 and consuming $c = g_4$. Column 2 shows the number of agents who use money (g_5) irrationally, column 3 shows the number of agents who use their production good irrationally (to purchase a higher storage-cost good), and column 4 shows the number of distinct ways of handling money (g_5) and the production good (g_2) actually represented in the economy.

6 When a High Transactions Cost Good is Money

It may be suspected that the genetic algorithm is simply picking out agents who are partial to goods with low storage cost. The next two simulations dispel this suspicion. In the first, we increase the frequency of g_4 -consumers from 20% to 44% of the total population, reduce the frequency of g_5 -consumers to 4% of the population. Since the probability of finding a buyer for g_4 is now significantly higher than for any other good, it is possible that g_4 serve as money for other agents, despite a storage cost that is higher than that of g_5 (for this simulation, the storage cost of g_4 was set to 3 and that of g_5 remained at 2).

The results are illustrated in Table 3. Using 2000 agents, the strategies stabilized after 935 generations, all agents accept g_4 as money (except g_3 -producers who consume g_1), and only g_1 -producers and/or consumers accept g_5 as money. This pattern of trade has an intuitive interpretation: storage costs are so high for g_1 -producers that they will accept g_5 in trade, despite its drawbacks; g_1 -consumers therefore accept g_5 because they can trade it with g_1 -producers (that g_3 -producers who consume g_1 do not accept g_4 is an anomaly that did not appear in other runs). This shows that both storage costs and frequency of use enter into the specification of money goods.

In our second simulation, we increase the frequency of g_5 -producers from

Agent Consumes	Agent Produces				
	g_1	g_2	g_3	g_4	g_5
g_1		g_1, g_4, g_5	g_1, g_5	g_1	g_1, g_3, g_4
g_2	g_2, g_4, g_5		g_2, g_4	g_2	g_2, g_4
g_3	g_3, g_4, g_5	g_3, g_4		g_3	g_3, g_4
g_4	g_2, g_3, g_4, g_5	g_3, g_4	g_4		g_2, g_3, g_4
g_5	g_3, g_4, g_5	g_3, g_4, g_5	g_2, g_4, g_5	g_5	

Table 3: **A High Transactions Cost Good is a (Nearly) Universal Medium of Exchange.** Shows the goods agents accept in trade for their production good, for all agent types, with storage costs (40,30,20,3,2), where the frequency of g_4 -consumers is 44%, and the frequency of g_5 -consumers is 4% of the population. All agents accept g_4 as money (except g_3 -producers who consume g_1), and only g_1 -producers and/or consumers accept g_5 as money.

20% to 68% of the total population, and we reduce the frequency of all other producer types equally, to 8% of the population each. There is now a glut of the low-storage-cost good g_5 , so we would expect agents to reject g_5 in trade, and move towards using g_4 as a money good. Again for this simulation, the storage cost of g_4 was set to 3 and that of g_5 remained at 2. The results are illustrated in Table 4. Using 4000 agents, the strategies stabilized after 760 generations, all agents accept g_4 as money, and only g_1 -producers and some g_1 -consumers accept g_5 as money. As in the previous simulation, storage costs are sufficiently high for g_1 -producers that they will accept g_5 in trade despite its oversupply, and some g_1 -consumers accept g_5 because they can trade it with g_1 -producers. This shows in addition to storage costs the frequency of production affects whether a good is accepted as money.

7 A Symbolic Money Equilibrium

In this section we investigate and simulate a Markov economy with *symbolic money*.

Definition 2 *A symbolic good is a good that is consumed by no agent. A Markov economy with symbolic good g is a Markov economy in which there is a good $g \in G$ that is not consumed. Symbolic money is a symbolic good that is accepted in trade by all agents. We say symbolic money is valuable if*

Agent Consumes	Agent Produces				
	g_1	g_2	g_3	g_4	g_5
g_1		g_1, g_4, g_5	g_1, g_5	g_1	g_1, g_3, g_4
g_2	g_2, g_4, g_5		g_2, g_4	g_2	g_2, g_4
g_3	g_3, g_4, g_5	g_3, g_4		g_3	g_3, g_4
g_4	g_2, g_3, g_4, g_5	g_3, g_4	g_4		g_2, g_3, g_4
g_5	g_3, g_4, g_5	g_3, g_4, g_5	g_2, g_4, g_5	g_5	

Table 4: **A Glut of the Low Transactions Cost Good Drives it Out of the Market.** Shows the goods agents accept in trade for their production good, for all agent types, with storage costs (40,30,20,3,2), where the frequency of g_5 -producers is 68%, and the frequency of all other producer types is 8% of the population. All agents accept g_4 as money (except g_3 -producers who consume g_1), and only g_1 -producers and/or consumers accept g_5 as money.

it is never part of an agent's Nash strategy to discard symbolic money.

Were it not for the fact that agents' inventories disappear when they die, valuable symbolic money could not exist in equilibrium. To see this we offer the following definition and theorem.

Definition 3 *A Markov economy with inheritance is a Markov economy in which each newly born agent inherits its deceased predecessor's inventory.*

Theorem 4 *Valuable symbolic money does not exist in equilibrium in a Markov economy with inheritance.*

Proof: Suppose m is valuable symbolic money in equilibrium. Since m is never consumed or discarded, and upon death an agent's inventory of m is passed on to its successor, the amount of m in circulation cannot decrease. Since the ergodic theorem implies that there is an equilibrium frequency of holding of m , if the probability of production of m is positive in a recurrent state, then the frequency of holding m must be unity, which implies that trade never occurs, so holding m is dominated by the no-trade strategy. This is a contradiction, which implies that m cannot be produced with positive probability in equilibrium. Consider an agent b who produces m and holds a nonempty inventory of k items in some recurrent state. Then b must face a positive expected probability of trading in each period, or else the no-trade strategy would dominate b 's strategy holding $k > 0$ items in inventory.

Thus a simple inductive argument shows that with positive probability b 's inventory will be empty in k periods, and hence b will produce a unit of m with positive probability in k periods. This shows that the assumption that $k > 0$ is in contradiction with the assumption that m is valuable symbolic money. Thus we must have $k = 0$, so an m -producer follows a no-trade strategy in any recurrent state. But by assumption all agents consider the economy to be in its ergodic state from birth, so no m -producer would have ever produced a unit of m , which is a contradiction, and hence the frequency of m in the economy is zero. This proves the theorem.

Without inheritance, symbolic money becomes possible in equilibrium, because new money can be created to take the place of money that disappears when its owner dies. It is clear, however, that if the death rate is low then very little trade can occur in an economy with valuable symbolic money, since m -producers must encounter mostly other agents holding nothing but m in inventory. Moreover, we would expect our simulations of a Markov economy with g_5 as a symbolic good to show that (a) when consumption utility is low, a no-trade equilibrium occurs; (b) when consumption utility is high, g_4 emerges as a universal medium of exchange and no agents accept g_5 in exchange; and (c) with sufficiently high consumption utility, if we start out with a high proportion of g_5 -acceptors, an equilibrium emerges in which g_5 is symbolic money and g_4 is also a medium of exchange .

These three expectations turn out to be correct. In our simulations of a Markov economy with symbolic good g_5 , we reduce the proportion of g_5 -producers to 8% of the total, and increase the others to 23% each, to lessen the problem of a g_5 -glut when there are no g_5 consumers. Since our simulations force agents to trade, the equivalent of a “no trade equilibrium” is for all agents to have negative utility in equilibrium, so they would prefer to hold empty inventories if allowed to do so. This is indeed what occurs when we set consumption utility to 150 when storage costs are (40,30,20,10,2) and production costs are 50 for each good.

If we increase consumption utility to 1000, the equilibrium described in (b) above emerges, as exhibited in Table 5. Despite its relative high storage cost, most agents accept g_4 as money (g_3 - and g_5 -producers who consume g_2 do not), and no agents accept the symbolic good g_5 as money.

Moreover if we increase consumption utility to 2000 and constrain all agents initially to accept g_5 in trade, after 1000 generations we attain a situation in which all agents (except g_5 -producers, of course) accept g_5 in trade, and all types of agents have positive levels of total wealth, indicating

Agent Consumes	Agent Produces				
	g_1	g_2	g_3	g_4	g_5
g_1		g_1, g_3, g_4	g_1, g_4	g_1	g_1, g_4
g_2	g_2, g_4		g_2	g_2	g_1, g_2, g_3
g_3	g_3, g_4	g_1, g_3, g_4		g_3	g_1, g_2, g_3, g_4
g_4	g_4	g_3, g_4	g_4		g_1, g_2, g_3, g_4

Table 5: **Low Transaction Cost Symbolic Good g_5 is Rejected as Money.** Shows the goods agents accept in trade for their production good, for all agent types, with storage costs (40,30,20,3,2), where the frequency of g_5 -producers is 8%, and the frequency of all other producer types is 23% of the population. All agents accept g_4 as money (except g_3 - and g_5 -producers who consume g_2), and no agents accept the symbolic good g_5 as money.

that their strategies dominate the no-trade strategy. There is, however, a very low level of trade in this case, with only 2.25% of the agents holding something other than g_5 in trade. Without constraining a high percentage of the initial population to accept g_5 in trade, this equilibrium would not have emerged. Indeed, Table 6 shows the result of constraining only 60% of the initial agents to accept g_5 in trade in this case. For the first forty generations there is a rapid increase in the acceptance of g_5 , but the agents who do so begin to be eliminated, and the frequency of holding of g_5 declines from 85% in round 40 to 21.3% by round 500, as most agents accept g_4 rather than g_5 as a medium of exchange.

8 A Markov Economy with *Fiat* Money

Definition 4 A fiat good in a Markov economy is a good m that is neither produced nor consumed, but may be traded, and upon the death of an agent, is passed on to the agent's successor. Fiat money in a Markov economy is a fiat good m that is accepted in trade by all agents.

The requirement that a fiat good be inherited is plausible and also necessary, since otherwise any such good would disappear from the economy in a finite number of periods.

Theorem 5 Consider a Markov economy with fiat money m such that each agent produces exactly one good, each good except m is produced by at least

Generation	Amount Offered for Sale				
	g_1	g_2	g_3	g_4	g_5
0	460	460	460	460	160
20	85	87	101	111	1616
40	50	67	60	123	1700
180	145	145	127	486	1097
300	295	210	221	659	615
400	250	280	316	699	455
500	254	303	311	706	426

Table 6: **Rejection of Symbolic Money when Initial Acceptance Rate is 60%**. For the first forty generations there is a rapid increase in the acceptance of g_5 , but from the 40th to the 500th generation the frequency of holding of g_5 declines from 85% to 21.3%. Most agents accept g_4 rather than g_5 as a medium of exchange (as indicated by the increase in holdings of g_4 over time).

one agent, and for any two goods $g, h \neq m$ there is a producer of g that consumes h . If the amount of m is sufficiently small, the storage cost of m is sufficiently low, and the utility of consumption is sufficiently high, there is an equilibrium in which each good $g \neq m$ trades for m with positive probability in each period.

Proof: We add three constraints to agent's pure trade strategies. First, we constrain each agent to hold a unit of some good in its inventory in each period. Second, we constrain agents to offer only one unit of a good in trade, (so in effect all trades are one-for-one). Third, we require all agents to accept a unit of m in exchange for some good $g \neq m$ currently held in inventory. By Theorem 1, there exists a Nash equilibrium for this constrained economy. To see that there is a positive probability of trade in this economy under the appropriate conditions, assume that there is only one unit of m in the economy, and its storage cost is less than that of any other good in the economy. Let a hold m in inventory. We first note that if a trades m for any good, a must trade m for consumption good c_a . Moreover, if the $u_a(c_a)$ is sufficiently high, a with an empty inventory has a positive present value (the proof of this mirrors the first part of the proof of Theorem 2). Hence a holding only m in inventory will trade m upon encountering an agent holding c_a , since otherwise a will hold m forever, which is dominated by trading m for c_a and then producing a unit of p_a . But with positive probability a 's

predecessor died in the previous period holding m , in which case a inherited m . It is clearly a strictly dominated strategy for the new a to add p_a to its inventory, since every agent who accepts p_a accepts m , and at most one unit can be traded in each period. Thus a holds only m in inventory. Then with positive probability a meets an agent b holding c_a in inventory (because with positive probability b is a c_a -producer who has just been born and hence must produce a unit of c_a for sale); as we have seen, a will offer to trade m for c_a , and b must accept this offer. Thus m is traded in each period with positive probability.

Now let us drop the constraints on a single agent a . We need only show that a will continue to obey the constraints. Since a has positive expected wealth when following the constrained strategy (assuming sufficiently high utility of consumption), a will not switch to a no-trade strategy, and hence will continue to hold a unit of some good in inventory in each period (the one-shot deviation principle is used here). Moreover a does not gain from offering more than one unit in trade, since an offer of more than one unit will not be accepted. Finally, since all agents who accept a good $g \neq m$ in trade for a good $h \neq m$ also accept m in trade for h , and since m has lower storage costs than any other good, a will continue to accept m in trade for any good in a 's inventory. Thus the constrained Nash equilibrium is an unconstrained Nash equilibrium, which proves Theorem 5.

Our simulations show that given the existence of a fiat good m , under reasonable conditions, 'self-fulfilling expectations' lead to m serving as a fiat money in equilibrium; i.e., if the initial stock of artificial life agents in a Markov economy who are constrained to accept m in trade is sufficiently large, then in equilibrium m is fiat money. We illustrate this in Table 7, which shows that if fiat good/agent ratio is 20% and we constrain the initial stock of agents so that 50% accept g_5 as fiat money, then in equilibrium g_5 serves as fiat money (in this case g_4 is money in equilibrium as well). Other simulations give the same result when the fiat good/agent ratio is a low and 10% and as high as 60%, and when as low as 25% of the original stock of agents is constrained to accept g_5 as fiat money.

If the initial stock of agents is randomly formed, however, the fiat good g_5 does not become fiat money. A typical simulation is illustrated in Table 8, where the fiat good/agent ratio is 30%. In equilibrium g_4 serves as money for all agents, but g_5 serves as fiat money only for producers and consumers of g_1 . Other simulations, with the fiat good/agent ratio varying from 10% to 50%, give virtually identical results. An obvious interpretation of this

Agent Consumes	Agent Produces			
	g_1	g_2	g_3	g_4
g_1		g_1, g_4, g_5	g_1, g_4, g_5	g_1, g_5
g_2	g_2, g_4, g_5		g_2, g_4, g_5	g_2, g_5
g_3	g_2, g_3, g_4, g_5	g_3, g_4, g_5		g_3, g_5
g_4	g_2, g_3, g_4, g_5	g_3, g_4, g_5	g_4, g_5	

Table 7: **A Self-Fulfilling Expectations Equilibrium with Fiat Money g_5 .** Shows the goods agents accept in trade for their production good, for all agent types, with storage costs (40,30,20,3,2), where g_5 is a fiat good, and 50% of the initial stock of agents are constrained to treat g_5 as fiat money. All agents accept both g_4 and g_5 as money.

results comes to mind. Since g_1 is the good with the highest storage costs, g_1 -producers quickly learn to exchange g_1 for g_5 just to reduce their storage costs. Thus g_1 -consumers find that it is worth while to hold g_5 because it is in demand by g_1 -producers. This increases the payoff to g_1 -producers of the ‘accept g_5 ’ strategy, which thus stabilizes in equilibrium.

Agent Consumes	Agent Produces			
	g_1	g_2	g_3	g_4
g_1		g_1, g_4, g_5	g_1, g_2, g_4, g_5	g_1
g_2	g_2, g_3, g_4, g_5		g_2, g_4	g_2
g_3	g_2, g_3, g_4, g_5	g_3, g_4		g_3
g_4	g_3, g_4	g_3, g_4	g_4	

Table 8: **An Evolutionary Equilibrium with Fiat Money g_5 .** Shows the goods agents accept in trade for their production good, for all agent types, with storage costs (40,30,20,3,2), where g_5 is a fiat good, and an initially random selection of agents. Only g_1 -producers and g_1 -consumers accept g_5 as fiat money.

9 Conclusion

The results presented in this paper could be strengthened in two obvious ways. First, our simulations generally converge to a unique ‘noisy’ equilibrium when started in a random state, while our formal models show that a

large number of Nash equilibria exist with widely differing qualitative properties. This difference in behavior results from the fact that the simulations use a replicator dynamic with random mutation and noise (in the form of random pairing of agents), while the formal models use the much weaker criterion embodied in the Nash equilibrium concept. The Nash equilibrium criterion should be replaced by a stronger dynamic criterion.²⁴ Second, we have made no attempt to model the process of relative price formation. We could easily have achieved more realism by restricting consideration to Nash equilibria in which relative prices are equal to relative costs of production. The resulting equilibria would appear more reasonable, but they in fact would not be in any substantive sense. Relative prices should flow from the equilibrium itself, and should reflect both costs of production and expected storage costs over the trading history of a good. While there are some possibly difficult technicalities involved, it should be possible to model relative price formation by having agents who meet for trade engage in a bargaining process of the Nash nash:coop or Rubinstein rubinstein82 variety.

²⁴Examples of dynamic modeling of this type include (?), (?), (?), and ?).

10 List of Symbols

A	Agents	P_a	Goods produced by a
		P^σ	Transition matrix with strategy σ
C_a	Goods consumed by a	s_{ag}	Cost to a of storing g for one period
		s_{ag}	Cost to a of storing i for one period
G	Goods for economy	pi^σ	ergodic probabilities for economy \mathbf{M}^σ
I_a	Admissible inventories for a	Σ_a	Pure strategies for a
I	Admissible inventories for the economy	Σ	Pure strategies for the economy
		Σ_a^*	Mixed strategies for a
		Σ^*	Mixed strategies for economy
κ_{ag}	Cost of production of g for a	u^a	Utility function of a
\mathbf{M}^σ	Markov chain when strategy is σ	Ω	Sample path space of Markov process of economy

11 Acknowledgments

I would like to thank Duncan Foley, Stephen Durlauf, participants of the Santa Fe Institute conference on Social Interactions and Aggregate Economic Activity for helpful suggestions in improving this paper, and the MacArthur Foundation for financial support.