Learning to Speculate:
Experiments with Artificial and Real Agents

John Duffy*
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260, USA

February 2000

Abstract

This paper employs an artificial agent–based computational approach to understanding and designing laboratory environments in which to test Kiyotaki and Wright's search model of money. The behavioral rules of the artificial agents are modeled on the basis of prior evidence from human subject experiments. Simulations of the artificial agent–based model are conducted in two new versions of the Kiyotaki–Wright environment and yield some testable predictions. These predictions are examined using data from new, human subject experiments. The results are encouraging and suggest that artificial agent–based modeling may be a useful device for both understanding and designing human subject experiments.

JEL classification: D83, C73, C90, E40.

Keywords: Learning, search, money, agent-based computational economics, experimental design.

*Tel: +1-412-648-1733; email: jduffy@pitt.edu; http://www.pitt.edu/~jduffy/jduffy.html.

1 Introduction

Monetary theorists have long struggled to understand why individuals choose to hold money when it is dominated in rate of return by other assets, for example, by interest bearing securities. Hicks, in his seminal (1935) paper, argued that understanding the rate of return dominance of money is "the central issue in the pure theory of money," and suggested that the necessary "frictions" that might help to explain this puzzle should be explicitly incorporated into the standard competitive general equilibrium model. Subsequently, researchers have proposed a number of departures from the standard model which allow money to be dominated in rate of return, for example, the inclusion of real money balances in the utility function, the requirement of cash in advance for purchases or the recognition of legal restrictions that force individuals to hold some amount of money (e.g. to pay taxes). However, the question remains as to how these devices arose in the first place. More fundamentally, one might ask how individuals came to accept or learned to adopt a convention in which the money asset is dominated in rate of return by other assets.

A promising environment in which to explore this important question is the decentralized, search–theoretic approach to monetary economics pioneered by Kiyotaki and Wright (1989) and extended by many others. Unlike the other modeling approaches mentioned above, search–friction models do not single out any special role to be played by a money good (as emphasized by Wallace (1998)) and allow for the endogenous determination of the good(s) that serve as media of exchange. Moreover, in the Kiyotaki–Wright (1989) model, it is possible for the goods that serve as media of exchange to be dominated in rate of return by other goods.

The essence of the Kiyotaki–Wright (1989) model is what Jevons termed the "absence of double coincidence of wants." In such environments, individuals must, of necessity, make investment decisions among competing goods since they cannot trade the good they produce directly for a good that yields them an immediate positive utility payoff. The game–theoretic aspect of the investment decisions embedded in this model is identical to that faced by investors in environments with network externalities when there is more than one convention or standard that might be adopted. In these situations, the relative expected discounted present values of the investment options facing any particular player at any point in time depend upon the current and future investment strategies adopted by all players. In certain circumstances, it may "pay" to make investments in goods that are relatively more costly to acquire today, as they provide a higher probability of a positive utility payoff tomorrow. It is in this sense that goods that are dominated in rate of return (are more costly to acquire) may nevertheless serve as media of exchange in the Kiyotaki–Wright environment. While monetary theorists have been attracted to the search–theoretic, dynamic game approach as a microfoundation for a general equilibrium model that explains the medium of exchange role of money in a natural way, an understanding of how a convention evolves when people are placed in a dynamic framework is of much greater general interest (See, e.g. Young (1998)).

Kiyotaki and Wright (1989) only provide a characterization of the equilibrium properties of their model. They do not attempt to describe the process by which an equilibrium is achieved, or equivalently, give an account of how one or more commodities might emerge as conventional media of exchange. Since agents almost certainly do not begin a process of social interaction with equilibrium beliefs but must adjust their strategies to their evolving historical experiences within a given trading regime, the relevant question is whether the comparative static predictions of the theory are actually informative of how play evolves when agents are placed in the Kiyotaki-Wright environment.

The evolution of behavior in the Kiyotaki–Wright (1989) environment has attracted the attention of many researchers, perhaps because the trade–offs that agents face in this environment appear simple enough that agents might be expected to learn the strategies that arise from solving a dynamic programming problem in a reasonable amount of time. The Nash equilibrium predictions of the model have been tested using populations of artificial, boundedly rational agents–the "agent–based computational approach"–by Marimon, McGrattan and Sargent (1990), Başçi (1999) and Staudinger (1998). The model's predictions have also been examined in a number of controlled laboratory experiments with paid human subjects ("real" agents) as reported in Brown (1996) and Duffy and Ochs (1999a, 1999b).

The major finding that emerges from all of these studies is the failure of both the artificial and real agents to engage in "speculative" behavior in certain though not all parameterizations of the model where speculative behavior constitutes a best response on the part of certain agent types.[1] 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. Thus, speculation is necessary for the

[1] In particular, Başçi (1999) and Staudinger (1998) find some parameterizations of the Kiyotaki–Wright model in which *artificial* agents do learn to play speculative strategies. The play of speculative strategies has not been observed human subject studies of the Kiyotaki–Wright environment described in this paper.

rate–of–return dominance feature of a medium of exchange.

While the "lack–of–speculation" finding is common to both the agent–based and human subject experiments, the design of these various experiments differ considerably. For example, all of the agent–based studies cited above involve many more artificial agents than are possible in a controlled laboratory environment with human subjects and these agent–based environments are run for many more periods than are possible in human subject experiments. Furthermore, all of the agent–based simulations allow for direct or indirect communication among agents (e.g. via their use of a genetic algorithm to update strategies) – a feature that is simply not present in human subject experiments. These differences make it difficult to directly compare the results from agent–based simulations with human subject experiments. Perhaps as a consequence, researchers have overlooked the possibility that agent–based simulations might serve as a tool for the design of human subject experiments, or that the results of human subject experiments might inform the modeling of agent behavior in artificial, agent–based economies. This paper takes a first step toward integration of the agent–based and human subject experimental approaches in the context of the much–studied Kiyotaki–Wright model.

We begin in the next section by briefly describing the Kiyotaki– Wright model. In section 3, we develop a simple, agent–based model and show that this model can match some of the features of the human subject experiments reported in Duffy and Ochs (1999a). The agent–based model is designed to be quite similar to the human subject experimental environment. Furthermore, the manner in which the artificial agents learn over time is based on findings from human subject experiments. In section 4, we use simulations of this agent–based model to predict what may happen in two previously unexplored versions of the Kiyotaki–Wright model that are designed to encourage greater speculative behavior by certain player types. In one new version of the model, agents who should speculate in the unique Nash equilibrium are given more frequent encounters with situations where playing the speculative strategy would result in higher expected utility. In the other new version of the model, two of the three agent types are constrained to playing deterministic strategies consistent with the Nash equilibrium where the best response of the other agent type is to speculate. The agent–based simulations predict different outcomes in these two model environments. The predictions of the agent–based simulations are then compared in section 5 with some human subject experiments conducted in the two new environments. The results are encouraging; the findings from the human subject experiments are roughly similar to those predicted by the agent–based simulations. Finally, section six offers some concluding remarks.

2 The Kiyotaki–Wright Environment

The discrete–time environment consists of a finite population of N players.[2] There are three player types, type $i = 1$, 2 or 3, with an equal number $(N/3)$ of each type. Type i desires to consume good i, but produces good $i+1$ modulo 3. Notice that with this specification, in the absence of trade, there can be no double coincidence of wants. There are no production costs however consumption of good i is a necessary prerequisite to production of good $i+1$ modulo 3. Goods are indivisible, and each agent has one unit of storage capacity in every period. The per period cost of storing a unit of good j is c_j and it is assumed that $0 < c_1 < c_2 < c_3$. This particular specification for storage costs, which we use throughout this paper, is referred to as "Model A" in Kiyotaki and Wright (1989).[3] All agents are assumed to receive the same utility from consumption, $u > c_3$, and to have a common discount factor, $\beta \in (0,1)$.

In every period, all N agents are randomly paired with one another. Each pair may engage in bilateral exchange if such an exchange is mutually agreeable. An exchange consists of a one–for–one swap of the goods the matched pair of players holds in storage when they enter the trading period. When an agent successfully trades for his consumption good i, he consumes that good, receiving utility u, and immediately produces a unit of his production good which he then stores until the next period, paying the per period storage cost for storing that good. Consequently, agent i is never storing a unit of good i; he enters each trading period with either good $i+1$ or $i+2$ in storage. We can therefore summarize the proportion of type i players who are storing their production good $i+1$ in period t by $p_i(t)$, with $1 - p_i(t)$ representing the proportion of type i players storing the other non–consumption good $i+2$, and let $p(t) = [p_1(t), p_2(t), p_3(t)]$ denote the vector of such proportions.

When an agent is unsuccessful in trading for his consumption good or does not face the opportunity to trade for his consumption good, his net payoff for the round is negative, corresponding to the storage cost of the good he holds in storage after any trading has occurred. Restrictions on the model's parameters imply that it is always a dominant strategy for player type i to offer to trade

[2] Kiyotaki and Wright (1989) assume an infinite continuum of players, but for our artificial and real subject experiments we must be content with a finite population approximation.

[3] An alternative specification, "Model B" (which is *not* isomorphic) reorders storage costs so that $0 < c_1 < c_3 < c_2$, or equivalently, changes the production specification so that type i produces good $i-1$.

whatever good he has in storage for his consumption good i as he receives a positive net payoff from such trades. The more interesting trading decisions arise in situations where a player type i with production good $i+1$ in storage, is randomly paired with a player storing good $i+2$. Denote the strategy of player i, who stores good $i+1$ and who meets good $i+2$ in period t by $s_i(t)$. Let $s_i(t) = 0$ if type i refuses to trade good $i+1$ for good $i+2$ and let $s_i(t) = 1$ if type i offers to trade good $i+1$ for good $i+2$. These strategies are assumed to be symmetric: if $s_i(t) = 0$ so that type i refuses to trade good $i+1$ for good $i+2$ at time t then $1 - s_i(t) = 1$, and type i offers to trade good $i+2$ for good $i+1$ at time t.

Kiyotaki and Wright assume that each agent type solves a dynamic programming problem in which they maximize the expected present value of discounted utility from consumption net of storage costs over an infinite horizon by choice of optimal trading strategies in all possible trading situations. They focus attention on pure strategies ($s_i(t) \in \{0,1\}\ \forall\ i,\ t$) and characterize the steady state equilibria that result from this optimization problem. A steady state, pure strategy Nash equilibrium consists of stationary vectors p and s such that the play of s results in the inventory distribution p and s_i is optimal for each player type i given p and s. For the case we study here, "Model A" where $0 < c_1 < c_2 < c_3$, Kiyotaki and Wright (1989) prove that there exists a unique, steady state pure strategy Nash equilibrium characterized by the strategy profile:

$$s = (s_1, s_2, s_3) = \begin{cases} (0,1,0) & \text{if } (c_3 - c_2) \;>\; [p_3 - (1-p_2)]/3\beta u = [1/6]\beta u, \\ (1,1,0) & \text{if } (c_3 - c_2) \;<\; [p_3 - (1-p_2)]/3\beta u = [(\sqrt{2}-1)/3]\beta u. \end{cases}$$

The equalities on the right hand side of the above expressions follow from the steady state vectors $p = (1, .5, 1)$ and $p = (.5\sqrt{2}, \sqrt{2}-1, 1)$, respectively. For $(c_3 - c_2) \in \left([(\sqrt{2}-1)/3]\beta u,\ [1/6]\beta u\right)$ there is no steady state pure strategy Nash equilibrium (though there do exist mixed strategy Nash equilibria).

Notice that, regardless of the parameter values chosen, the unique pure strategy Nash equilibrium calls for type 2 players to always offer to trade their high storage cost good 3 for the less–costly–to store good 1 ($s_2 = 1$), and for type 3 players to always refuse to offer to trade their low storage cost good 1 for the more costly–to–store good 2 ($s_3 = 0$). These trading strategies, in which players trade high storage cost goods for lower storage cost goods or refuse to trade low storage cost goods for higher storage cost goods, are referred to as *fundamental* trading strategies, since fundamental factors, i.e. storage cost considerations, govern the trading decisions.

By contrast, type 1 players' optimal pure strategy depends on the choice of model parameter values. Under certain parameterizations the optimal pure strategy for type 1 players, $s_1 = 0$, is the fundamental strategy in which type 1 players refuse to trade their lower cost production good 2 for the more costly to store good 3. However, there exist parameterizations for which type 1 player's optimal pure strategy is the *speculative* strategy, $s_1 = 1$, of trading their production good 2 for the more costly to store good 3. This speculative strategy is optimal whenever the difference in costs from storing good 3 rather than good 2 ($c_3 - c_2$) is less than the discounted expected utility benefit of storing good 3 rather than good 2, which is given by $[p_3 - (1 - p_2)]/3\beta u$. Intuitively, type 1 players are better off speculating in good 3 (rather than good 2) if storing good 3 makes it more likely that they will be able to successfully trade for their desired consumption good – good 1 – and the additional likelihood of this event more than outweighs the additional cost of storing good 3 rather than good 2.

Our focus in this paper will be on parameterizations of the model that give rise to the unique strategy profile $s = (1, 1, 0)$. In such versions of the model, there are two *media of exchange* defined as goods that are accepted in trade but not desired for consumption purposes. Good 1 fulfills this role, as type 2 players trade for good 1, but desire good 2. Good 3 also fulfills this role as type 1 players trade for good 3 but desire good 1. Since good 3 is the most costly–to–store good, yet is accepted in trade by certain player types it serves as an example of an endogenously determined medium of exchange that is dominated in rate of return.

3 Artificial agent behavior in the Kiyotaki–Wright model

3.1 Prior analyses

Kiyotaki and Wright (1989) only characterize the steady state pure strategy Nash equilibria of their model. They do not provide an analysis of the process by which a population of inexperienced or boundedly rational agents might learn over time to coordinate on this equilibrium. A number of researchers have sought to explore this follow–up question in the context of the Kiyotaki–Wright model using either agent–based computer simulations or controlled laboratory studies with paid human subjects, but not both.

The earliest study, by Marimon, McGrattan and Sargent (1990), used a version of Holland's (1986) classifier system to model the behavior of a population of 150 artificial agents in various versions of the Kiyotaki–Wright (1989) model (with 50 agents per type). Each artificial agent, of a given player type, has a collection of state–contingent "if–then" rules for both trading and consumption decisions. Each rule is assigned a strength according to its contribution to past realized utility using Holland's "bucket brigade" credit assignment procedure; this algorithm distributes strengths to the sequence of all rules that contributed to the utility gain or loss, rather than to the single rule that actually induced a change in the agent's utility. Periodically, new if–then rules are generated using a genetic algorithm; these rules replace poorly performing rules. The classifier rules with the highest strengths are the ones chosen to represent each agent's state–contingent behavior. Marimon et al. (1990) report that when the model environment is parameterized so that the unique strategy profile across the three agent types is $s = (0, 1, 0)$, so that all types play fundamental strategies in the steady state, the classifier system has no trouble achieving the unique pure strategy "fundamental" equilibrium. However, when the model is parameterized so that the unique strategy profile is $s = (1, 1, 0)$, the artificial type 1 agents fail to learn to play the required speculative strategy of trading their production good 2 for the more costly–to–store good 3 at the end of 1,000 periods.

More recently, Başçi (1999) has conducted further classifier system simulations in the Kiyotaki–Wright (1989) model using populations of size 60 (20 agents per type). Başçi considers two main modifications to the Kiyotaki–Wright environment in an effort to induce a greater percentage of type 1 players to adopt the speculative strategy: 1) reducing the differential in storage cost between goods 3 and 2, $(c_3 - c_2)$ and 2) allowing a certain fraction of agents to choose strategies according to their social i.e. population–wide values, which Başçi refers to as "imitation." Başçi finds that neither modification by itself results in a significant increase in the speed of convergence to the speculative strategy profile, $s = (1, 1, 0)$, but that the combination of these two modifications does enable the classifier system to achieve convergence on this steady state with a high frequency, (as great as 89 % of all runs) at the end 1,000 iterations. Staudinger (1998) uses a pure genetic algorithm implementation (as opposed to a classifier system) to model the behavior of a population of 150 agents (50 per type) in the Kiyotaki–Wright model. Like Başçi, Staudinger varies the storage cost differential, $c_3 - c_2$, and finds that for small enough differentials, the genetic algorithm is always able to achieve the speculative strategy profile $s = (1, 1, 0)$, by the end of 5,000 generations (periods).

These studies, while informative, are not easily compared with laboratory studies, e.g. by Brown (1996) or Duffy and Ochs (1999a, 1999b), where relatively smaller population sizes of 18–30 agents (6, 8 or 10 players of each type) are used in each experimental session. Moreover, the artificial agent simulations are conducted for many more periods (rounds) than are possible in controlled laboratory studies, where concerns about subject boredom limit the amount of time that can be spent in the laboratory. Finally, we note that the various algorithms used to model learning in the artificial agent simulations are not easily implementable or feasible in laboratory environments. For example, all three studies discussed above make use of a genetic algorithm. The interpretation of the reproduction and crossover operations of the genetic algorithms is that agents somehow communicate with one another about which strategies are relatively more fit (yield higher payoffs), and in the "crossover" operation, (the main operator of a genetic algorithm) parts of these highly fit strategies (e.g. if–then rules) may be combined with other strategy components to achieve new classifiers or strategies. In laboratory experimental environments it is difficult to control communication among agents in the same manner as prescribed by a genetic algorithm or in a manner that leads to a clear understanding of what is going on. For these reasons, communication among agents is typically not allowed in human subject experiments though it is possible to provide subjects with aggregate, population–wide information.

Clearly what is needed is to design artificial agent environments in a manner that allows them to be tested in laboratory environments with paid human subjects. While this approach will certainly limit the kinds of behavior and information processing that are possible in artificial agent simulations, the reward is a more disciplined and empirically–grounded approach to the modeling of artificial agent behavior as well as the use of such agent–based models to design better experimental environments for human subject experiments. We now turn our attention to developing such an artificial agent model for the Kiyotaki–Wright environment.

3.2 An artificial agent model based on experimental findings

The artificial agent environment is very similar to the human subject environment examined by Duffy and Ochs (1999a).[4] In particular, there is a total of just 24 or 18 artificial agents. These agents are divided up equally among the three player types, so that there are 8 or 6 players of each type. Each artificial agent simulation (or "run") is conducted in the same manner as a single experimental session with human subjects; the simulation consists of a number of games, and each game consists of a number of rounds. While each human subject experiment involved

[4]For the complete details of this environment including the experimental design and instructions, the reader is referred to Duffy and Ochs (1999a).

a different group of 18 or 24 inexperienced subjects, in each artificial agent simulation we use a different random seed for the pseudo–random number generator that decides the outcome of the probabilistic decisions made by agents (as discussed below).[5]

In every round of a game, players are randomly paired with one another. They then decide whether to offer to trade the goods they have in storage. As in the human subject experiment, there is a .90 probability that the game continues from one round to the next, and a .10 probability that the current round will be the last round of the game. This stochastic end to a game was designed to implement a discount factor, β of 0.90, while simultaneously implementing the stationarity that is associated with the model's infinite horizon. By contrast, in the artificial agent world, stationarity issues are of little concern as long as agents are not conditioning their decisions explicitly on time. However, to facilitate comparison with the experimental data, we have adopted the same probabilistic procedure for ending a game in the artificial agent simulations.[6] We also follow the Duffy–Ochs experimental protocol and end each artificial agent simulation after approximately 100 rounds of play, or an average of about 10 games.

As our aim is to develop an agent–based model to design further experiments, it seems reasonable to let earlier experimental findings aid in the design of the agent–based model. We followed the design and findings from the human subject experiments in constructing and simplifying the decision rules followed by the artificial agents in our agent–based model. First, as in the human subject experiments, but in contrast to the classifier system used by Marimon et al. (1990), our artificial agents do not make consumption decisions; if a type i player succeeds in trading for good i that agent immediately consumes good i and produces a unit of his production good, $i+2$ modulo 3. Second, the experimental evidence suggests that players nearly always (97% of the time) offer to trade for their consumption good whenever they are paired with another agent who is storing this good. This is not surprising given that such behavior is a dominant strategy; agents receive a positive net payoff only when they consume. In our artificial agent simulations, we therefore chose to make the decision of a type i agent who is facing the opportunity to trade for good i *deterministic*: the type i agent always offers to trade for good i. Finally, we note that the random matching technology allows two agents to meet who both have the same type of good in storage.

[5]It is also the case that a different random seed is used to generate the sequence of random pairings of agents in both the artificial simulations and human subject experiments.

[6]One feature of this design that may be important is that at the start of each new game, subjects begin with some good in storage that may differ from the good they held in the last round of the just completed game.

Since storage costs are the same across agent types, there is no reason for such a pair to engage in trade; there is absolutely no effect on either agent's utility if a trade occurs or fails to occur. Therefore, in such cases, we impose the deterministic rule that the two players both refuse to engage in trade. While this last restriction would seem to be straightforward we note that in the simulations of Marimon et al. (1990), no such restriction is made.

A type i player's critical decision is whether or not to trade his production good $i+1$ for the other possible non–consumption good $i+2$. Consider a player type i who is storing good $i+1$ at the end of period t. Suppose also that the random draw is such that the game "continues" into period $t+1$ and our player meets another player who is storing good i. As noted above, our artificial player type i will always offer to trade whatever good he has in storage for good i. Suppose that trade is mutually agreed upon; the player type i trades for and consumes good i earning utility u. Let the type i player's utility gain from storing good $i+1$ at time t be denoted by $\gamma_{i+1} = -c_{i+1} + \beta u$. Similarly, if player i is storing the other non–consumption good $i+2$ at the end of period t and successfully trades this good for the desired consumption good in the following period $t+1$ his gain from storing good $i+2$ at time t is denoted $\gamma_{i+2} = -c_{i+2} + \beta u$.

On the other hand, suppose the type i player is storing good $i+1$ when he meets good i in trade, but is *unsuccessful* in trading for good i. The only explanation for this failure is a refusal to trade on the part of the other player. Since strategies in our learning model are assumed to be symmetric, we infer that this other player would have offered to trade good i to our player type i if our type i player was storing his other non–consumption good, good $i+2$. Accordingly, the *opportunity cost* to our player type i of storing good $i+1$ and *failing* to trade for good i in the following round is γ_{i+2}, the utility gain he would have experienced had he held the other good $i+2$ in storage.

We use these utility gains and opportunity costs to determine the probability that each player of type i offers to trade good $i+1$ for good $i+2$. The modeling of this probabilistic decision is based on the experimental findings of Duffy and Ochs (1999a), in particular, by their logit regression model specification for individual subject behavior, which they found to provide a reasonably good fit to the experimental data. Let the net payoff in period t to individual agent j, of type i, from storing good $i+1$ be denoted by

$$\nu^j_{i+1}(t) = \sum_{\tau=1}^{t-1} I^s(\tau)\gamma_{i+1} - \sum_{\tau=1}^{t-1} I^f(\tau)\gamma_{i+2},$$

where $I^s(\tau)$ is an indicator function that is equal to 1 if player j was storing good $i+1$ in period τ and was *successful* in trading it for good i, and is equal to zero otherwise. Similarly $I^f(\tau)$ is an indicator function that is equal to 1 if player j was storing good $i+1$ in period i and *failed* to trade it for good i, and is equal to zero otherwise. The net payoff in period t to individual j of type i from storing good $i+2$, denoted $\nu^j_{i+2}(t)$, is similarly defined. Notice that these net payoffs will differ over time for each individual agent j due to both the random matching technology and differences in the goods held in inventory among agents of the same type. Note also that the payoff to storing good $i+1$ or good $i+2$ gets reinforced by the actual net utility gain in the event that the agent is successful in trading either of these two goods for good i. However, when the agent fails to trade either good $i+1$ or good $i+2$ for good i, the payoff to storing that good is negative and equal to the foregone net utility gain the agent would have received had he held the other good in storage. This "hypothetical reinforcement" approach differs from a strict reinforcement learning approach, e.g Roth and Erev (1995), but is consistent with other learning approaches, for example stochastic fictitious play (as described e.g. in Fudenberg and Levine (1998)).

The relative benefit to individual agent j of type i from storing good $i+1$ rather than good $i+2$ is defined by the difference, $x^j_i(t) = \nu^j_{i+1}(t) - \nu^j_{i+2}(t)$. Using this difference, we define the probability that agent j of type i plays strategy $s_i = 0$ by:[7]

$$\Pr[s^j_i(t) = 0] = \frac{e^{x^j_i(t)}}{1 + e^{x^j_i(t)}}$$

Recall that the strategy $s_i = 0$ is the strategy of refusing to trade the type i agent's produced good $i+1$ for good $i+2$. The probability that agent j of type i plays strategy $s_i = 1$ (trades good $i+1$ for $i+2$) is then given by $1 - \Pr[s^j_i(t) = 0]$.

This logistic specification was chosen because it provided a good fit to the behavior of human subjects facing the decision to trade good $i+1$ for good $i+2$ as reported in Duffy and Ochs (1999a). Similar specifications for probabilistic choice models have become popular among researchers seeking to understand polarized outcomes at an aggregate level that arise from localized interaction, (e.g. the statistical mechanics approach to understanding socioeconomic heterogeneity advocated by Durlauf (1997) and others).

In the human subject experiments conducted by Duffy and Ochs (1999a) as well as those

[7]More generally, we could add parameter values to this probabilistic strategy model, for example we could multiply the difference $x^j_i(t)$ by a parameter for scaling purposes and we could also add a constant term to the difference. For simplicity, we avoid such complications.

described below, subjects were able to see on their computer screens the total number of points they had received from all trading rounds within a game as well as the effect of the decisions made in each trading round on these point totals. Subjects earned a positive number of points equal to u only when they successfully traded for their consumption good. They lost points every round according to the fixed storage cost (c) of the good they held in storage at the end of the trading round. At the end of an experimental session, subjects' end–of–game point totals from one game, chosen at random from all the games played in that session, were converted into a probability of winning a \$10 prize (which was in addition to a \$10 participation payment). Each additional point that subjects earned in the game chosen increased their probability of winning the \$10 prize by the same amount. Since the game chosen to determine this probability was not known in advance, the "real" agents had the incentive to obtain as many points as possible in every round of every game. Analogously, the artificial agents' reinforcements, as described above, depend on the decisions they make in every round of every game.

3.3 Comparison of Simulation Results with Experimental Results

In all of the simulation exercises reported in this paper we use the parameterization of the model that is given in Table 1. This parameterization of the model was also used by Duffy and Ochs (1999a) in one of their main experimental treatments, thereby facilitating a direct comparison.[8]

Table 1: Model Parameters

u	1.00
c_1	0.01
c_2	0.04
c_3	0.09
β	0.90

These parameter choices imply that $s = (1, 1, 0)$ is the unique Nash equilibrium strategy profile. Hence, type 1 players are called upon to play speculative trading strategies, trading the good they produce, good 2, for the more costly to store good 3 whenever they have the opportunity, whereas type 2 and 3 players are called upon to play fundamental strategies. The resulting steady state

[8] The values for u, c_1, c_2 and c_3 in Table 1 differ only in magnitude from the values used by Duffy and Ochs (1999a) by a scale factor of 100. There are no consequences of this rescaling for the predictions of the model.

vector for the proportion of each type storing their production good is $p = (.71, .59, 1.0)$.

Duffy and Ochs'(1999a) experiments were conducted under two different schemes for the initialization of goods across player types at the start of each new game. Here we adopt the initialization scheme in which the initial distribution of goods $i+1$ and $i+2$ over players of type i was made as close as possible (with a finite population) to the steady state proportions of these two goods; this is the most severe test of the stability of the steady state. We will compare our simulation results using this initialization scheme with human subject experimental results that used this same initialization scheme.

Our first set of 5 artificial agent simulations were conducted with 24 artificial agents per simulation run. These agents were divided up equally across the three player types (8 players of each type). These simulations are compared with the results of 5 human subject experimental sessions with 24 or 18 subjects as reported in Duffy and Ochs (1999a). In each of the human subject experiments there were equal numbers of each of the three player types (8 or 6).

Some aggregate statistics from the Duffy-Ochs (DO) human subject experiments and the artificial agent (A) simulation runs are provided in Table 2 below. This table reports the frequency with which each player type $i = 1, 2, 3$ played strategy $s_i = 1$ over each half of an experimental session or artificial simulation approximately 50 rounds per half).

Table 2: Offer Frequencies Over Each Half of 5 Sessions with Real or Artificial Agents

Real Session Number	Number of Subjects	Type 1 Offers 2 for 3		Type 2 Offers 3 for 1		Type 3 Offers 1 for 2	
		First Half	Second Half	First Half	Second Half	First Half	Second Half
DO1	24	.13	.18	.98	.97	.29	.29
DO2	24	.38	.65	.95	.95	.17	.14
DO3	24	.48	.57	.96	1.00	.13	.14
DO4	18	.08	.24	.92	.98	.12	.02
DO5	18	.06	.32	.93	.97	.25	.18
All	108	.23	.37	.95	.96	.20	.16
Artificial Session No.	Number of Subjects	Type 1 Offers 2 for 3		Type 2 Offers 3 for 1		Type 3 Offers 1 for 2	
		First Half	Second Half	First Half	Second Half	First Half	Second Half
A1	24	.06	.15	.73	1.00	.37	.07
A2	24	.23	.31	.88	.98	.20	.07
A3	24	.33	.50	.78	.98	.15	.00
A4	24	.18	.42	.81	1.00	.17	.00
A5	24	.10	.18	.67	.98	.23	.07
All	120	.19	.32	.77	.99	.22	.04

The simulated sessions yield aggregate offer frequencies that are similar to those of the experimental sessions. In particular, we find that in the first half of each human subject session, the average frequency with which type 1 players offer to trade good 2 for good 3 across 5 sessions is 19%; in the simulations the comparable frequency is 23%. This average offer frequency increases to 32% in the second half of the human subject experiments as compared with an average offer frequency of 37% in the simulated sessions. Recall that in this parameterization, type 1 players are expected to play the speculative strategy of always offering to trade good 2 for good 3. It appears our agent–based model does a reasonable job capturing the reluctance of type 1 players to engage in such speculative behavior. Furthermore, the aggregate offer frequency statistics for the other two player types as generated by our agent–based learning model also compare favorably with those of the human subject experiments, though there are some differences. Indeed, it appears that the artificial type 2 players take longer on average to learn to trade their production good 3 for the lower cost good 1 as compared with experimental subjects. Also, speculative behavior by the artificial type 3 agents dies out in the second half of sessions at a faster rate than in the experimental data. Nevertheless, qualitatively, the "fit" of the artificial agent simulation statistics to those of from the experimental data appears to be quite good.

In addition to tracking the aggregate behavior of agents in the Kiyotaki–Wright model, our agent–based learning model is also able to capture *individual* agent behavior. For example, one might ask whether the offer frequencies reported for type 1 players in Table 2 suggest that all type 1 players are playing mixed strategies, or that some type 1 players are "fundamentalists" playing the fundamental strategy of refusing to trade good 2 for good 3, while other type 1 players are "speculators" offering to trade good 2 for the more costly to store good 3. In both the human subject experiment and in the artificial agent environment the answer is the same: some type 1 players are strict fundamentalists, never experimenting with the speculative strategy and there are also some type 1 players who learn to play the speculative strategy early on, and tend to stick with this strategy over much of the remainder of the session (run). There does not appear to be much evidence of any artificial or real player playing a mixed strategy by the end of a session or run.

[Insert Figures 1a–1b here.]

Figure 1a shows the trading behavior over all rounds of a session for each of the eight type 1 players in a representative human subject experimental session and Figure 1b does the same for a representative agent–based simulation run. In these figures a "1" indicates a decision by each type

1 player to trade good 2 for good 3, and "-1" indicates a refusal to offer to make such a trade; a "0" indicates that the type 1 player with good 2 was not facing the opportunity to trade for good 3. Each type 1 player's action choices, separated by vertical bars, are shown in sequence from the first to the last round of the session. The figure clearly reveals that there is heterogeneity among both the real and artificial type 1 players, who appear to be mainly playing one or the other of the two possible pure strategies.

The logistic model used in the artificial agent simulations can account for this kind of polarized division of agents between fundamentalists and speculators, as early differences in reinforcements move individual agents in the direction of one of the two pure strategies. These differences in the reinforcements that agents attach to the two pure strategies need not disappear over time, as agents' trading incentives are in turn altered by the heterogenity in strategic behavior. Indeed, early adoptions of pure strategies appear to get reinforced and are frequently adhered to for the duration of a run. Hence it is initial experience with a certain pure strategy, which differs across agents due to the random matching process as well as the probabilistic trading rules, that largely accounts for the observed heterogeneity across individual agents in the artificial agent simulations.

4 Two New Treatments

While Duffy and Ochs' (1999a) findings called into question the plausibility of Kiyotaki and Wright's speculative pure strategy Nash equilibrium, Duffy and Ochs were careful to note that there might exist different ways of "framing" or presenting the problem to subjects that would make the benefits of the speculative strategy relatively more transparent, thereby increasing the likelihood that type 1 players adopted this strategy. Duffy and Ochs did assess the robustness of their results to changes in the parameterization of the model and to different information conditions, but did not consider more substantial changes in the experimental design as budget and time limitations prevented such an endeavor. Nevertheless, it is important to consider whether there exist changes in the experimental design that might lead to greater adoption of the speculative strategy in order to properly assess the generality of Duffy and Ochs' findings. Rather than pursue such a search with further human subject experiments, a more cost–effective approach might be to use simulations of agent–based models to test experimental designs prior to implementing those designs in the laboratory with paid human subjects.

Since the artificial agent–based model we developed in the last section appears to replicate some of the aggregate and individual findings of the Duffy–Ochs (1999a) experiment, it seems reasonable to use simulations of this model to forecast the consequences of changes in the experimental design on human subject behavior. In considering alternative designs, the aim was to find a design that might hasten convergence toward the unique pure strategy Nash equilibrium where type 1 players are called on to play speculative strategies, within the length of time (approximately 100 rounds) of a human subject session. Two new designs were considered, and we now turn our attention to describing these designs and the predictions of our agent–based learning model.

4.1 The Role of Experience: An Unequal Distribution of Players Across Types

The first new experimental design is based on Wright (1995), who modified the original Kiyotaki–Wright environment to allow an unequal distribution of players across types. In particular, the total population of players remains fixed at size N, but the fraction of type 2 players is reduced below 1/3 while the fraction of type 3 players is increased above 1/3. The fraction of type 1 players is assumed to remain constant at 1/3. This unequal distribution of players across types implies that type 1 players will have more frequent encounters with type 3 players relative to the case in which players are equally distributed across types and thus type 1 players will gain greater experience with situations in which the speculative strategy might yield more frequent opportunities to trade for and consume good 1 (therefore reinforcing speculative play). Wright (1995) found that this distribution of players across types was the limiting distribution obtained when agents were allowed to choose their type and storage costs were such that $c_1 < c_2 < c_3$ (as assumed here).[9]

The parameterization used in the agent–based simulation involved a population size of $N = 18$ agents, as we planned to conduct the human subject experiments using just 18 subjects. Consistent with prior experiments, one–third of the players (a total of 6) were assigned the role of type 1 agents. The remaining 12 players were either type 2 or type 3 players. We chose to have 4 type 2 players (2/9) and 8 type 3 players (4/9) as this distribution of players across types was a close as possible with a finite population to the limiting distribution of players across types that emerges from Wright's simulations where agents can choose their type. In all other respects, the model environment is unchanged. In particular, the model parameterization is the same as in the earlier

[9] In particular, Wright allowed the fraction of agents assigned to the role of type 2 or 3 players to evolve according to a simple replicator dynamic. He reports the limiting distribution of players over types from many simulations of this replicator dynamic.

set of experiments and it is easily shown that the unique Nash equilibrium strategy vector remains $s = (1, 1, 0)$.

We performed 5 simulation runs using the artificial agent economy. The only difference between this artificial environment and the one used in the simulations reported in Table 2 was in the number of agents assigned to play the role of each player type. The aggregate statistics on the frequency with which each type offers to trade good $i+1$ for good $i+2$ are reported in Table 3.

Table 3: Offer Frequencies Over Each Half of 5 Artificial Agent Sessions with Unequal Division of Players Across Types.

Session Number	Number of Subjects	Type 1 Offers 2 for 3		Type 2 Offers 3 for 1		Type 3 Offers 1 for 2	
		First Half	Second Half	First Half	Second Half	First Half	Second Half
A6	18	.53	.70	.84	1.00	.12	.00
A7	18	.28	.65	.85	.96	.15	.00
A8	18	.44	.47	.66	.97	.15	.00
A9	18	.12	.65	.80	1.00	.23	.00
A10	18	.59	.92	.90	1.00	.20	.00
All	90	.40	.67	.81	.98	.17	.00

A comparison of the simulation results in Table 3 with those given in Table 2 reveals that the treatment in which players are unequally distributed across types does appear to have an effect on the aggregate offer frequencies. In particular, we see that by the second half of each session a majority of type 1 players are playing speculative trading strategies, trading good 2 for good 3, and all type 3 players are adhering to the fundamental strategy of refusing to trade good 1 for good 2. However, we still do not see perfect coordination on the predicted Nash equilibrium strategy vector in any session by the end of the approximately 100 rounds of play per simulation run.

4.2 Eliminating Noise: Automating the Decisions of Type 2 and Type 3 Players

In an effort to further hasten convergence to the unique Nash equilibrium in which type 1 players speculate while type 2 and 3 players play fundamental strategies, we considered a second treatment in which we eliminate learning behavior on the part of type 2 and 3 players. Instead, the decisions of type 2 and 3 players are automated – that is, they are made independent of these player's personal histories of play. We modeled type 2 and 3 players as always playing according to their predicted (i.e. steady state) fundamental trading strategies. In particular, we suppose that type 2 players always offer to trade their production good 3 for the less costly to store good 1 and that type 3 players always refuse to trade their production good 1 for the more costly to store good 2. By eliminating "noisy" trading behavior on the part of type 2 and 3 players we hoped to reduce the time it takes type 1 players to learn to play the speculative strategy of trading good 2 for good 3, and thus hasten convergence to the unique Nash equilibrium of the model environment. In essence, this treatment transforms the experiment from a group decision making problem to an individual decision one, in the sense that only type 1 players have to learn the equilibrium strategy to play.

We again performed 5 simulation runs using our artificial agent economy. In this new treatment, the behavior of type 2 and 3 players was deterministic, as described above, while the behavior of type 1 players was governed by the same history–dependent probabilistic rule used in the previous simulations. Here, we returned to the case of having equal numbers of the three agent types. In particular, we assumed a population size of 24 agents, with 8 players assigned to play the role of each player type, as we planned to conduct experiments with 8 type 1 subjects. The aggregate statistics on the frequency with which each type offers to trade good $i+1$ for good $i+2$ are reported in Table 4.

Table 4: Offer Frequencies Over Each Half of 5 Artificial Agent Sessions with Deterministic Type 2 and 3 Players

Session	Number of	Type 1 Offers 2 for 3		Type 2 Offers 3 for 1		Type 3 Offers 1 for 2	
Number	Subjects	First Half	Second Half	First Half	Second Half	First Half	Second Half
A11	24	.57	.73	1.00	1.00	.00	.00
A12	24	.54	.62	1.00	1.00	.00	.00
A13	24	.60	.74	1.00	1.00	.00	.00
A14	24	.71	1.00	1.00	1.00	.00	.00
A15	24	.70	.74	1.00	1.00	.00	.00
All	120	.62	.73	1.00	1.00	.00	.00

The interesting offer frequencies in Table 4 are those for type 1 players, as in this treatment, type 2 and 3 players are always playing according to the fundamental trading strategy. We see that the elimination of noisy trading behavior on the part of type 2 and 3 players has an effect that is qualitatively similar to our treatment with an unequal division of players across types; the frequency with which type 1 players play the speculative strategy of trading good 2 for good 3 is higher than in the treatment where type 2 and 3 players' strategies are probabilistic (compare the offer frequencies in Table 4 with those in Table 2). In fact, the simulated type 1 player offer frequencies in Table 4 are slightly higher, on average over the first and second halves of a session, than those in Table 3 where players were unequally divided up among the three player types. In one of the five simulations reported in Table 4, A14, the artificial agent economy actually achieves convergence to the steady state Nash equilibrium of the model within the alloted number of trading rounds (again approximately 100 rounds per session).

4.3 Discussion

Taken together, the results of these two sets of agent–based simulations suggest that it may be difficult, though not impossible, to achieve convergence to the unique "speculative" Nash equilibrium within the short amount of time that is allowed in economic decision making experiments. Since the second new treatment, where type 2 and 3 players are automated, yielded somewhat faster progress toward convergence, on average, than the first new treatment, it would be reasonable to choose this design as the one to pursue for further experimental testing with paid human subjects. However, as our aim is to evaluate the predictive power of this agent–based model for purposes of experimental design, we chose to conduct experimental tests with paid human subjects using *both* of the new experimental treatments. The agent–based simulations predict that 1) both of the new treatments lead to higher frequencies of the play of speculative strategies by type 1 players relative to the baseline, Duffy–Ochs (1999a) treatment and 2) convergence is slightly faster on average in the treatment with automated type 2 and 3 players as compared to the treatment with an unequal division of players across types.

5 Experiments with real agents

The human subject experiments were conducted in the same manner as in Duffy and Ochs (1999a), as briefly described above. However, as the two new treatments studied in this paper involve modifications to the Kiyotaki–Wright environment, some changes were required in the design of the human subject experiments relative to the design used by Duffy and Ochs.

The number of human subjects in each experimental session was made equal to number of artificial agents in each run of the agent–based learning model, as indicated in Tables 3 and 4. As in Duffy and Ochs, the human subjects had no prior experience with the Kiyotaki-Wright model. Each subject participated in only one session. Subjects were randomly assigned to play the role of a single player type for the entire session. In the first treatment, with an unequal division of players across types, a total of 18 subjects were used in each of two experimental sessions. As in the artificial agent simulations, six of these human subjects were assigned to play the role of type 1 players, 4 subjects were assigned to play the role of type 2 players, and 8 subjects were assigned to play the role of type 3 players. In the second treatment with automated type 2 and 3 players, 8 human subjects were recruited per session (two sessions were conducted) and all 8 subjects were assigned to play the role of type 1 players. In these sessions there were 8 automated type 2 players and 8 automated type 3 players for a total of 24 "subjects" per session. The human subjects in this second treatment were informed that the type 2 and 3 players they encountered through random matches were automated and the subjects were also informed of the strategies these automated players were playing, e.g. which goods the automated players would offer and refuse to trade for.[10] The type 2 and 3 players were programmed to always play the fundamental strategies predicted by the unique, pure strategy Nash equilibrium.

In all other respects, the experimental design was the same as in Duffy and Ochs (1999a). Subjects began each game with 100 points and either their production good, or their other non–consumption good in storage. The initial distribution of goods over player types was the same one used in the artificial agent simulations. This distribution was made as close as possible to the unique steady state distribution of goods over player types. The model parameters used in the human subject experiments were the same as those used in the artificial agent experiments (Table 1) but multiplied by a scale factor of 100. Thus, for instance, a type 1 player who successfully traded for good 1 received 100 points less the cost of 4 points for storing a unit of his production good, good 2, for a net gain of 96 points. Subjects lost points in accordance with the per round storage cost of the good they were holding at the end of any round in which they failed to trade for their consumption good, or traded for some good other than their consumption good. Thus, as noted earlier, subjects received or lost the same point amounts in each period as the artificial agents (adjusting for the scale factor). Following the Duffy–Ochs design, at the end of each human subject session, one of the games that subjects played in the session was chosen at random. The subjects' point totals from that game were converted into a probability of winning a $10 prize (which was in addition to a $10 participation payment) The probability of winning this prize was an increasing,

[10]Note that the artificial agents do not make use of such information. This information concerning the strategies played by the automated type 2 and 3 players was provided to the human subjects playing the role of type 1 players to convince these subjects that they were playing automated players. Revealing information about the strategic behavior of type 2 and 3 players is also in keeping with the idea of eliminating noisy information about type 2 and 3 player's decisions.

linear function of the number of points that subjects earned in the game chosen. Since subjects did not know in advance the game that would be chosen, they had the same incentive as the artificial agents to earn as many points as they could in every round of every game.

As in prior experiments and the artificial agent simulations each game ended randomly, with a constant probability of termination of .10 from one round of a game to the next. Subjects in each session played around 100 rounds, or approximately 10 games per session. The subjects were paid a $10 participation fee. In addition, one of the games the subjects played was chosen at random. Subjects' point totals from that game were converted into a probability of winning an additional $10 prize. The more points subjects earned in each game, the greater was their probability of winning this additional prize if the game was chosen at the end. Thus subjects had the incentive to get as many points as possible in each round of each game.

5.1 Unequal Distribution of Players Across Types

Table 5 reports results from the first two "real" human subject experimental sessions, R1 and R2, of the treatment with an unequal division of players across types. This table gives the aggregate frequencies with which each player type offered to trade his production good for the good he neither produces nor consumes over the first and second half of each session (approximately 50 rounds each) and over both sessions combined. We see that, consistent with the predictions of the artificial agent simulations, there is some increase in the frequency with which type 1 players in both sessions play speculative strategies in the first and second half of both sessions as compared with earlier experimental findings reported in Table 2. Moreover, we find that the increase in the frequency of speculative play by type 1 players, though slight, is similar in magnitude to the average frequencies of speculative play across the five simulated sessions involving artificial agents.

Table 5: Offer Frequencies Over Each Half of 2 Human Subject Sessions with Unequal Division of Players Across Types.

Session	Number of	Type 1 Offers 2 for 3		Type 2 Offers 3 for 1		Type 3 Offers 1 for 2	
Number	Subjects	First Half	Second Half	First Half	Second Half	First Half	Second Half
R1	18	.33	.48	.98	1.00	.29	.13
R2	18	.60	.79	.96	.97	.21	.20
Both	36	.46	.64	.97	.98	.25	.16

However, some caution seems warranted in interpreting these findings as support for the notion that this new treatment hastens convergence toward the unique speculative equilibrium. There is simply not enough evidence yet (with just two sessions) to conclude that the treatment with an unequal division of players across types yields offer frequencies for type 1 players that are significantly different from those reported in Table 2 for an equal division of players across types. The results reported in Table 5, however, are encouraging.

5.2 Automated Type 2 and 3 Players

The aggregate offer frequencies from the real human subject sessions, R3 and R4, where the decisions of type 2 and 3 players were automated, are reported in Table 6. Despite the fact that type 2 and 3 players were always playing according to their steady state strategies, and type 1 players were made aware of this fact, we still do not find successful coordination on the speculative strategy by all type 1 players, even by the second half of the two sessions. Nevertheless, the average type 1 player offer frequencies over these two sessions are higher than in both the treatment with an unequal distribution of players across types, and in the baseline Duffy–Ochs (1999a) experiment. This finding is again consistent with the predictions of the agent–based model. Indeed, these frequencies of speculative play by type 1 players are the highest ever reported for this parameterization of the model.

Table 6: Offer Frequencies Over Each Half of 2 Human Subject Sessions with Automated Type 2 & 3 Players

Session Number	Number of Subjects	Type 1 Offers 2 for 3		Type 2 Offers 3 for 1		Type 3 Offers 1 for 2	
		First Half	Second Half	First Half	Second Half	First Half	Second Half
R3	24	0.84	0.83	1.00	1.00	0.00	0.00
R4	24	0.52	0.53	1.00	1.00	0.00	0.00
Both	48	0.69	0.71	1.00	1.00	0.00	0.00

Notice however, that there is some variance in the aggregate frequencies reported in these two sessions. The frequencies with which type 1 players play the speculative strategy in session R4 are somewhat lower than in session R3. Furthermore, the aggregate frequencies of speculative play in one session of the treatment with an unequal distribution of players across types, session R2 (as reported in Table 5), is higher than in session R4. (A similar variance in aggregate offer frequencies was found in the artificial agent simulations). Clearly, further experiments with paid human subjects are needed to determine whether there is a significant difference between the two new treatments. However, these preliminary experimental results, as reported in Tables 5 and 6 are generally supportive of the predictions of the artificial agent model concerning the effect of the new treatments on the adoption of speculative strategies.

6 Conclusion

Understanding why individuals demand money even though it is dominated in rate of return by other stores of value is a problem that has challenged monetary theorists. The search–theoretic approach to money, as pioneered by Kiyotaki and Wright (1989), offers a possible resolution to this modeling problem. Under certain parameterizations of this model, it is optimal for certain agent types to trade a relatively less costly–to–store good for a relatively more costly–to–store good if these agents have the rational expectation that they will be able to trade the more costly–to–store good relatively more quickly for the good they desire to consume. Previous experiments with real and artificial agents have shown that agents have difficulty learning to adopt such "speculative" trading strategies. In this paper we have used artificial, agent–based simulations to explore some simple modifications to the basic Kiyotaki–Wright environment that might make it easier for subjects to learn to play speculative trading strategies. The behavior of the artificial agents was constructed on the basis of observed characteristics of human subject behavior in a prior experimental study of the Kiyotaki–Wright model. The two modifications to the Kiyotaki–Wright model were: 1) an unequal distribution of players across types and 2) automating the decisions of all player types who are not called upon to play speculative strategies in equilibrium. The artificial agent simulations suggest that both of these modifications serve to increase the speed with which players learn to adopt speculative strategies, by comparison with earlier experimental Kiyotaki–Wright environments, and that the second new treatment has the relatively greater impact on the speed of learning. Preliminary experiments with human subjects seem to confirm the predictions of these artificial agent simulations.

One of the purposes of conducting this exercise was to suggest ways in which the agent–based computational modeling approach might be usefully combined with the methodology of controlled laboratory experiments with human subjects. I have suggested how prior experimental findings might serve to inform the modeling of artificial agent behavior, and how the restrictions of a laboratory environment need to be imposed on artificial agent–based models before comparisons can be made between the two approaches. Once such restrictions are in place, agent–based models can be a very useful instrument in both understanding and designing economic–decision making experiments, as suggested by the findings reported in this paper.

Acknowledgments

This paper builds on earlier work with Jack Ochs, and I thank him for many helpful discussions. I also thank an anonymous referee and Leigh Tesfatsion for their comments and suggestions, and Jeff Kinross for excellent research assistance. Research support from the National Science Foundation is gratefully acknowledged. The C++ source code for the artificial agent simulations as well as the data from the human subject experiments reported in this paper are available from the author upon request.

References

Başçi, E., 1999, Learning by imitation, Journal of Economic Dynamics and Control 23, 1569–1585.

Brown, P.M., 1996, Experimental evidence on money as a medium of exchange, Journal of Economic Dynamics and Control 20, 583–600.

Duffy, J. and J. Ochs, 1999a, Emergence of money as a medium of exchange: An experimental study, American Economic Review 89, 847–877.

Duffy, J. and J. Ochs, 1999b, Fiat money as a medium of exchange: Experimental evidence, working paper, University of Pittsburgh.

Durlauf, S.N., 1997, Statistical mechanics approaches to socioeconomic behavior, in: W.B. Arthur et al., eds., The economy as an evolving complex system II, (Addison–Wesley, Redwood City, CA).

Fudenberg, D. and D.K. Levine, 1998, The theory of learning in games, (MIT Press, Cambridge, MA).

Hicks, J.R., 1935, A suggestion for simplifying the theory of money, Economica 2, 1–19.

Holland, J.H., 1986, Escaping brittleness: The possibilities of general–purpose learning algorithms applied to parallel rule–based systems, in R.S. Michalski et al., eds., Machine learning: An artificial intelligence approach, (Morgan Kaufmann, Los Altos, CA).

Kiyotaki, N. and R. Wright, 1989, On money as a medium of exchange, Journal of Political Economy 97, 924–954.

Marimon, R., E.R. McGrattan and T.J. Sargent, 1990, Money as a medium of exchange in an economy with artificially intelligent agents, Journal of Economic Dynamics and Control 14, 329–373.

Roth, A.E. and I. Erev, 1995, Learning in extensive–form games: Experimental data and simple dynamic models in the intermediate term, Games and Economic Behavior 8, 164–212.

Staudinger, S., 1998, Money as a medium of exchange: An analysis with genetic algorithms, working paper, Technical University, Vienna.

Young, H.P., 1998, Individual strategy and social structure: An evolutionary theory of institutions, (Princeton University Press, Princeton, NJ).

Wallace, N., 1998, A dictum for monetary theory, Federal Reserve Bank of Minneapolis Quarterly Review 22, 20–26.

Wright, R., 1995, Search, evolution and money, Journal of Economic Dynamics and Control 19, 181–206.

Figure 1a: Action Choices by Type 1 Players Who Face Opportunity to Trade Good 2 for Good 3, Human Subject D02 (8 Type 1 Subjects)

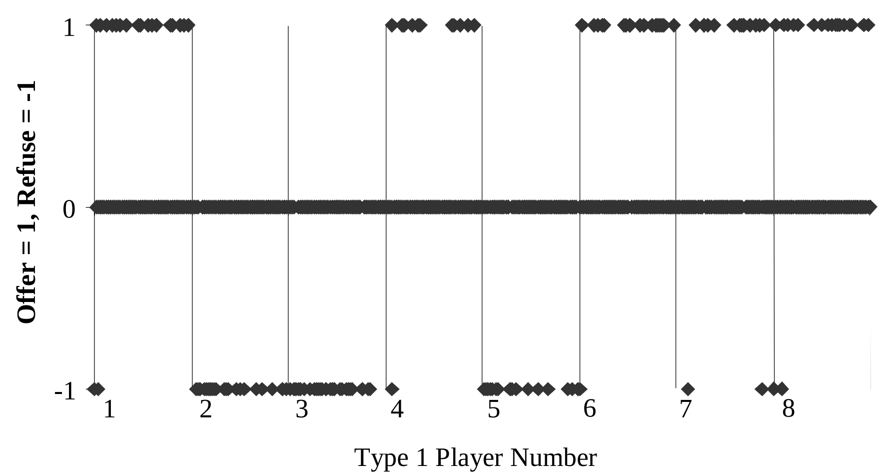

Figure 1b: Action Choices by Type 1 Players Who Face Opportunity to Trade Good 2 for Good 3, Artificial Agent A2 (8 Type 1 Subjects)