Emergence of Money as a Medium of Exchange: An Experimental Study

By JOHN DUFFY AND JACK OCHS*

This paper reports findings from an experiment that implements a search-theoretic model of money as a medium of exchange. The question examined is whether subjects learn to adopt the same commodities as media of exchange that the model predicts will be used in equilibrium. We report that subjects have a strong tendency to play "fundamental" rather than "speculative" strategies even in environments where speculative strategies yield higher payoffs. We examine some possible motivations for subjects' behavior and conclude that subjects are mainly motivated by past payoff experience as opposed to the marketability considerations that the theory emphasizes. (*JEL* D83, E40)

One of the most important conventions a society develops is the acceptance of at least one object as a "medium of exchange." An object becomes a medium of exchange when many agents who have no interest in that object for their own consumption or use in production nevertheless accept the object in trade with the rational expectation that they will be able to trade it for goods which are of intrinsic value to themselves. In this sense, a conventional medium of exchange must be supported by a Nash equilibrium. We all know when some commodity has acquired the status of a medium of exchange. The challenge to monetary theorists is to identify the factors that determine which object(s) will acquire this status.

Nobuhiro Kiyotaki and Randall Wright (1989) provide a model of an economy with individuals of several types where different commodities may emerge as media of exchange.[1] In this model each player type is defined by a pair of goods—the good a type consumes and the good a type produces. No type produces the good it desires to consume and there does not exist any pair of types in which each type produces what the other type wishes to consume. Each agent has one unit of storage capacity. If an individual has his own consumption good in storage, it is immediately consumed and replaced by that individual's production good. Therefore, each agent has in storage either a unit of his production good or a unit of a good other than his consumption good. At the beginning of each period, agents are randomly paired. If they mutually agree, the pair may exchange inventories. Therefore, consumption is only possible if at least some trades take place in which at least one agent accepts a good which that agent does not consume. Such goods are considered *media of exchange* and are identified via the equilibrium trading strategies of the agents and by the patterns of trade and inventory holdings induced by those strategies.

In the Kiyotaki-Wright model, different goods have different storage costs. Therefore,

* Department of Economics, 4S01 Forbes Quad, University of Pittsburgh, Pittsburgh, PA 15260. Financial support for this project was provided by the National Science Foundation under Grant No. SBR95-07968. We thank Gabriele Camera, Joe Ritter, Neil Wallace, Randy Wright, two anonymous referees, and participants at a number of seminars for their comments and suggestions. We are grateful to Chris Briem and Nick Meggos for research assistance and to Robert Conticello, Sridhar Komandur, and Jie Xing for programming assistance. The raw data from the experimental sessions reported in this paper are available from the authors upon request.

[1] Kiyotaki and Wright (1989, 1991, 1993) and S. Rao Aiyagari and Neil Wallace (1992) also consider the possibility that *fiat objects* serve as media of exchange. In this paper, we focus on the emergence of valued *commodities* as media of exchange. An experimental investigation of the emergence of fiat objects as media of exchange is provided in Duffy and Ochs (1999).

whenever an agent is paired with someone who does not have that agent's consumption good in storage the agent must decide which of two goods that differ in their storage costs it would be better to hold. The number of periods that will pass before an agent meets someone who both has that agent's consumption good in storage and is willing to accept in exchange the good the agent has in storage is itself dependent on which good the agent has in storage. Therefore, the decision regarding which good to hold in storage involves weighing the differences in the storage costs of the goods against the differences in their "marketability." Kiyotaki and Wright provide a characterization of the equilibrium properties of two versions of their model. They provide conditions under which various goods will serve as media of exchange in equilibrium. These conditions depend on the distribution of storage costs, the utility value of consumption, and the discount factor. What makes the Kiyotaki-Wright model important for monetary theory is that it can account for the existence of multiple media of exchange, some of which are dominated in rate of return (the financial asset analog of storage cost) by other possible media.

The interesting equilibria arise when parameter values imply that one or more types adopt "speculative" trading strategies. A *speculative* strategy is one in which an agent accepts in trade a good with a higher storage cost than the good that agent is currently holding. The agent speculates that by incurring a higher storage cost in this period he will experience a shorter wait before being able to effect a trade for his consumption good. In those parameter conditions in which the equilibria imply that high storage cost goods are used as media of exchange such "speculative" beliefs are, of course, self-fulfilling. Speculative trading strategies stand in contrast to "fundamental" trading strategies. A *fundamental* strategy is one in which an agent will accept in trade only those goods that have a lower storage cost than the good the agent is currently storing. This strategy is considered fundamental in that the agent considers only the fundamental factor—the storage cost of a good—when deciding whether or not to trade for that good. Like speculative strategies, fundamental strategies are self-fulfilling only in those environments where parameter conditions imply that low storage cost goods are used as media of exchange.

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 emerge as conventional media of exchange. As with other equilibrium analyses, this leaves open the question of whether or not the equilibria are likely to be achieved by agents who, almost certainly, do not start a process of social interaction with equilibrium beliefs but who must adjust their strategies to their evolving historical experiences within a given trading regime. Such a question can only be satisfactorily addressed by actually bringing people together in an environment specified by the model and observing their behavior.

I. An Experimental Study

This paper reports results from an experiment we have conducted that implements several parameterizations of the Kiyotaki and Wright (1989) environment. Our aim in conducting this experiment is to address the question of whether the equilibrium predictions of the Kiyotaki-Wright model are robust to the dynamics created by out-of-equilibrium play.

Our experimental design uses two principal treatment variables. One treatment variable is the distribution of production goods assigned to types. The two possible distributions, described below, are referred to as model A and model B following Kiyotaki and Wright (1989). A second treatment variable is the payoff value, or utility value, of an agent's consumption good. Under model A, the utility value of consumption is set either at a level that is consistent with fundamental behavior by all player types in equilibrium or with speculative behavior on the part of some types in equilibrium. Under model B, the utility value is set either to a value that is consistent with fundamental behavior by all player types in equilibrium or to a value that is consistent with either fundamental or speculative behavior by various types in equilibrium. In addition to these treatment variables, we also test whether manipulations in the form of public knowledge or the distribution of initial inventory holdings have any systematic effect on behavior.

We find, contrary to the theory, that in model A the aggregate relative frequency with which all types choose the fundamental strategy is unaffected by the choice of parameter values or by the manner in which goods are initially distributed over types. When model B is played, the aggregate relative frequency of selection of the fundamental strategy is also essentially the same, regardless of whether the parameters imply a unique equilibrium in which only fundamental strategies are predicted, or whether the parameters imply a multiplicity of equilibria, where selection of speculative strategies might be expected.

The theoretical equilibrium is predicated upon strategy selection that is conditioned on the differences in holding costs relative to the differences in the marketability of different commodities. Therefore, we investigate the way in which individuals respond to the kinds of information available to them. Our investigation suggests that in selecting a strategy, available information on the relative marketability of different goods does not influence a player's strategy choice. Given this behavior, it is not surprising that strategy selection is unresponsive to parameter variations that change the theoretical equilibrium.

The robustness of the Kiyotaki-Wright model's predictions to out-of-equilibrium behavior has been previously examined by Ramon Marimon et al. (1990), who used a version of John H. Holland's (1986) classifier system to model how a population of artificially intelligent players might interact in the Kiyotaki-Wright environment. They found that the artificial agents had no difficulty learning to play fundamental strategies. However, in environments where the unique pure strategy equilibrium involves speculative trading strategies, the artificially intelligent agents failed to adopt these speculative trading strategies, and continued to play fundamental strategies. Marimon et al. (1990 pp. 361–2) comment on the failure of their algorithm to learn to play speculative strategies: "In the limit, the artificially intelligent agents should behave as long-run average payoff maximizers (...) [but] the behavior of our artificially intelligent agents can be very myopic at the beginning (...) [and] this early myopia might have a perverse effect (...)." That is, play of a speculative strategy involves accepting a lower current payoff (higher storage cost) in exchange for a higher expected future payoff. Without sufficient experience of high payoffs from speculation, the speculative strategies get selected against early on and "die out."

Motivated by Marimon et al.'s findings, Paul M. Brown (1996) conducted an experimental test of the Kiyotaki-Wright model's predictions using human subjects. These experiments involved only model version A and were parameterized in a way that would call for speculative behavior by type 1 players in the unique pure strategy equilibrium. Brown finds that less than one-half of player type 1s chose speculative strategies when it was optimal for them to do so. He also found that one-sixth of player type 3s played the speculative strategy even though the unique pure strategy equilibrium of model A requires that all player type 3s play the fundamental strategy.

While we obtain aggregate results that are similar to the results obtained by Marimon et al. and Brown, our experimental investigation differs from these earlier studies in several respects. First, as mentioned above, we use an experimental design where we vary both the version of the model and the parameterizations of these different versions. By contrast, Brown considered only a single version of the Kiyotaki-Wright model (model A) and a single parameterization of this version of the model. Second, we are interested not only in aggregate behavior, which is the main focus of the Brown and Marimon et al. studies, but also in the motivations for why individual subjects choose particular strategies. We therefore develop a simple model of individual learning behavior which we investigate using data collected from the experiment. Finally, there are important differences between Brown's experimental design and our own. Brown's experimental design was motivated by the work of Marimon et al. and was therefore set up so as to replicate the environment in which the artificially intelligent agents played the game. However, this implementation leaves out certain features of the Kiyotaki-Wright model that would seem to be important to players' ability to achieve coordination on a Nash equilibrium, especially one that involves speculative trading strategies. For instance, in the Marimon et al. and Brown environments, no effort is made to implement discounting, as the theory requires. Furthermore, in the Brown experiment, no attempt is made to control for subjects' risk attitudes. Most importantly, in the Brown and Marimon et al. environments, individual players only have knowledge of their own past history of play. The absence of individual knowledge of population-wide information may limit the ability of players in a disequilibrium setting to achieve coordination on a Nash equilibrium.

This point has been emphasized in related work by Rajiv Sethi (1996). He shows that all of the pure strategy equilibria identified in the Kiyotaki and Wright (1989) model (including equilibria where some player types are called on to play speculative strategies) are asymptotically stable with respect to a disequilibrium, evolutionary dynamic that encompasses a large class of selection dynamics. This evolutionary dynamic can only operate in environments where agents have access to population-wide information that enables them to assess the relative performance of the strategies they are playing.

An important feature of our experimental design is that we have made an effort to provide subjects with this kind of population-wide information. In particular, we always provide subjects with knowledge of sufficient statistics for the population distribution of strategies that have been played, and attempt to make this fact common knowledge by reading the instructions aloud to all players at the beginning of each session. In some sessions, we provide subjects with additional, population-wide strategic information. We have also implemented discounting and have attempted to control for subjects' risk attitudes in contrast to the earlier studies. We therefore view our experimental design as being a closer approximation to the Kiyotaki-Wright environment.

The rest of the paper is organized as follows. Section II describes the versions of the Kiyotaki-Wright environment that we study and provides the conditions under which the various pure strategy equilibria exist. Section III provides the details of our experimental design. In Section IV we consider some aggregate characteristics of the experimental data and compare these results with the theoretical predictions of the model. In Section V, we present results from a regression analysis that is intended to identify the types of information that subjects used to inform their strategy choices. Finally, Section VI provides a summary of the major findings and some concluding remarks.

II. The Kiyotaki-Wright Environment

In the Kiyotaki-Wright model, a population of N players is divided up equally into one of three types.[2] There are also three indivisible goods. A player of type $i = 1, 2, 3$ has access to a technology that produces good $j = 1, 2, 3$ where $i \neq j$. In model A type $i = 1, 2, 3$ produces good $j = 2, 3, 1$, respectively, while in model B type $i = 1, 2, 3$ produces good $j = 3, 1, 2$, respectively. In both of these models, players of type i receive positive utility only from consumption of good i. Thus, type 1 players desire good 1, type 2 players desire good 2, and type 3 players desire good 3. The production possibilities of models A and B are designed to motivate indirect exchange of goods between agents of different types. In addition to a production technology, agents have access to a storage technology that allows each agent to store one unit of any good in every period. Storage is costly, and the costs differ according to which good is stored. Denoting the cost of storing good j by c_j, the storage costs that all players face are such that $c_3 > c_2 > c_1 > 0$.

In every period or trading round, all players are randomly paired and each pair may choose to trade the goods they have in storage. Trades must be mutually agreed upon by both players and involve one-for-one swaps of goods in inventory. If one or both members of a pair do not want to trade, no trade occurs and the two players exit the trading round storing the same good they held at the beginning of the round. If both players agree to trade then the goods they have in inventory are swapped.

All three player types $i = 1, 2, 3$ receive the

[2] More precisely, Kiyotaki and Wright assume a *continuum* of players of unit mass, but for our purposes, it will be necessary to assume a finite population size. The model can be generalized to allow for more than three player types as in Aiyagari and Wallace (1991), or to allow agents to choose their own type as in Wright (1995).

same level of utility, $u > c_3$ from consumption of good i. They also incur some disutility in every round corresponding to the storage cost of the good they are holding at the end of the round. When a player of type i successfully trades for his consumption good he immediately consumes that good and produces a new unit of his production good. He therefore earns positive net utility for the round in the amount $u - c_{i*}$, where $i*$ denotes type i's production good. Thus, no player ever stores the good he desires to consume; the good held in storage will either be the player's production good or the good the player neither produces nor desires to consume. The two other possible trading outcomes are that a trade is made but the player does not receive his consumption good in trade, or trade is not mutually agreeable, so no trade occurs. In either of these two cases, the player's net utility for the round is negative and is given by $-c_j$ where j is the good the player had in storage as of the end of the trading round.

Players maximize the discounted expected value of net utility over an infinite horizon by choice of optimal "trading strategies." The discount factor, $\beta \in (0, 1)$, is assumed to be common across types. A trading strategy is a rule that determines whether player type i holding good j offers to trade good j for some other good k. The players' problem is solved by applying standard dynamic programming techniques and by searching over all possible stationary pure strategies.[3] The solution to this problem can be characterized by the different proportions of agents storing the three goods in a steady-state equilibrium. Denote by p_{ij}, the proportion of type i agents storing good j and note that for agents of type i, $p_{ii} = 0$ by design. It follows that $p_{ij} = 1 - p_{ik}$ where $i \neq j \neq k$. Thus, an equilibrium inventory distribution is completely characterized by three steady-state proportions—one for each player type.

The model parameterizations that we consider in this study are provided in Table 1 using

TABLE 1—MODEL PARAMETER VALUES USED

Parameter	Model A value	Model B value
u	20 or 100	20 or 500
β	0.90	0.90
c_1	1	1
c_2	4	4
c_3	9	9

the notation discussed above.[4] Note that the only parameter values that differ within and across model treatments are the values for u, the utility value of consumption. Using these different sets of parameter values we will now characterize the conditions under which the various pure strategy equilibria of the Kiyotaki-Wright model exist.

Consider first, the case of model A where types $i = 1, 2, 3$ produce goods $j = 2, 3, 1$, respectively. In this version of the model, it is the actions of type 1 players alone that determine whether a pure strategy equilibrium is labeled as "fundamental" or "speculative." A player type 1 "plays fundamental" if he refuses to trade his production good 2 for the higher storage cost good 3. In model A, a type 1 player's steady-state best response is to play this fundamental strategy whenever:[5]

$$(c_3 - c_2) > (p_{31} - p_{21})\beta u/3.$$

This inequality says that the difference in storage cost between goods 3 and 2 exceeds the discounted expected utility benefit to player type 1s from storing good 3 rather than good 2. Using our parameterizations for model A, as reported in Table 1, we can rewrite the above inequality as:

$$(1) \qquad p_{31} - p_{21} < \begin{cases} \frac{5}{6} & \text{if } u = 20, \\ \frac{1}{6} & \text{if } u = 100. \end{cases}$$

If these inequalities hold, then type 1 players'

[3] Timothy J. Kehoe et al. (1993) extend the Kiyotaki and Wright (1989) model to allow for mixed strategy and nonstationary equilibria. In this paper, we follow Kiyotaki and Wright (1989) and focus on stationary, pure strategy equilibria.

[4] In two experimental sessions (discussed in Section IV, subsection B) we considered parameterizations that were slightly different from the values given in Table 1.

[5] See Kiyotaki and Wright (1989) for the derivation of this inequality and others that follow.

steady-state best response is to always play the fundamental strategy. For the other two types of players, types 2 and 3, the fundamental pure strategy of only trading for lower storage cost goods (or refusing to trade for higher storage cost goods) is always a steady-state best response for all valid parameterizations of the model. Of course, all three player types should always offer to trade for their consumption good as well, as they earn the highest possible payoff from doing so.

If all three player types always adhere to fundamental trading strategies, the pure strategy equilibrium is referred to as a "fundamental equilibrium." In the fundamental equilibrium of model A, type 1 players always refuse to trade their production good 2 for the higher storage cost good 3, type 2 players always offer to trade their production good 3 for the lower storage cost good 1, and type 3 players always refuse to trade their production good 1 for the higher storage cost good 2. Since good 1, the least costly-to-store good, is accepted in trade by type 2 players who do not desire to consume it, good 1 is regarded as a *medium of exchange* in this equilibrium. The steady-state inventory distribution in this pure strategy equilibrium is given by $p_{12} = 1.00$, $p_{21} = 0.50$, and $p_{31} = 1.00$, implying that $p_{31} - p_{21} = 0.50$. It is easily verified that with this inventory distribution, inequality (1) is *self-fulfilling* when $u = 20$; in the case where $u = 100$, this inequality will be violated, indicating that type 1 players' steady-state best response is to pursue a *speculative* strategy.

In the "speculative equilibrium" of model A, type 1 players play the speculative strategy in which they always offer to trade their production good 2 for the higher storage cost good 3. Player types 2 and 3 continue to adhere to fundamental strategies as before. Therefore, both goods 1 and 3 serve as media of exchange in this speculative equilibrium of model A. The resulting steady-state inventory distribution is given by $p_{12} = 0.71$, $p_{21} = 0.59$, and $p_{31} = 1.00$, implying that $p_{31} - p_{21} = 0.41$. Again, it is easily verified that with this inventory distribution, inequality (1) is violated when $u = 100$. It follows that the speculative pure strategy equilibrium, where player type 1s always offer to trade good 2 for good 3, will be self-fulfilling when $u = 100$.

Consider next the case of model B where types $i = 1, 2, 3$ produce goods $j = 3, 1, 2$, respectively. In this version of the Kiyotaki-Wright model, a pure strategy fundamental equilibrium where all three types play fundamental strategies *always exists,* for all valid parameterizations of the model. In this fundamental equilibrium, type 1 players offer to trade their production good 3 for the lower storage cost good 2, type 2 players refuse to trade their production good 1 for the higher storage cost good 3, and type 3 players offer to trade their production good 2 for the lower storage cost good 1. Thus, both goods 1 and 2 serve as media of media of exchange in the fundamental equilibrium of model B. The fundamental pure strategy equilibrium inventory distribution in model B is given by $p_{12} = 0.29$, $p_{23} = 0.00$, and $p_{32} = 0.41$.

Speculative play is also possible in model B. In contrast to model A, it is player types 2 and 3 who play speculative strategies in the "speculative equilibrium" of model B, while player type 1s continue to play the fundamental strategy. Another difference in model B is that when the speculative pure strategy equilibrium exists, it *coexists with* the fundamental pure strategy equilibrium, so that within a certain region of the parameter space, there are *two* possible pure strategy Nash equilibria that agents may choose to coordinate on. The conditions under which the speculative equilibrium coexists with the fundamental equilibrium are:

$$(c_3 - c_1) < (p_{32} - p_{12})\beta u/3,$$

$$\text{and } (c_2 - c_1) < p_{23}\beta u/3.$$

If both of these inequalities are satisfied, then there exists an equilibrium in which type 3 players speculate by refusing to trade their production good 2 for the lower storage cost good 1, and in response, type 2 players speculate by trading their production good 1 for the higher storage cost good 3. Type 1 players continue to adhere to the fundamental strategy in which they always trade their production good 3 for the lower storage cost good 2. In this equilibrium, goods 2 and 3 serve as media of exchange.

Given our parameterization of model B we can rewrite the above inequalities as:

$$(2)\quad (p_{32} - p_{12}) > \tfrac{4}{3} \quad \text{and} \quad p_{23} > \tfrac{1}{2} \quad \text{if } u = 20,$$

$$(3)\quad (p_{32} - p_{12}) > \tfrac{4}{75} \quad \text{and} \quad p_{23} > \tfrac{1}{50} \quad \text{if } u = 500.$$

The speculative pure strategy equilibrium inventory distribution in model B is given by $p_{12} = 0.59$, $p_{23} = 0.29$, and $p_{32} = 1.00$. Thus, in the speculative equilibrium, we have that $(p_{32} - p_{12}) = 0.41$. It is easily verified, using the inequalities given in (2–3), that this speculative equilibrium distribution will be *self-fulfilling* only in the parameterization where $u = 500$. Thus, the parameterization where $u = 500$ is one that is consistent with a Nash equilibrium where player type 2s and 3s play speculative strategies, and type 1 players play the fundamental strategy. Of course, both parameterizations, $u = 20$ and $u = 500$, are also consistent with the Nash equilibrium where all player types play fundamental strategies.

III. Experimental Design

In implementing the Kiyotaki-Wright environment in the lab we have had to make some simplifying assumptions. To begin with, in our experimental implementation we are constrained to using a *finite* population of players as opposed to the model's infinite continuum. Another important simplification is our use of a finite horizon rather than the infinite horizon of the theory. We will address the infinite horizon issue later in our discussion.

We have chosen our finite population of N players to be relatively large, at least by comparison with other experimental studies. However, the size of our experimental populations was constrained by the number of computer workstations available in our experimental laboratory (30), and by the rate at which subjects showed up for the experiment. In addition, our population sizes had to be multiples of 6 due to the model's requirement that there be equal numbers of each of the three types and due to the use of pairwise matching. For these reasons, each of our experimental sessions involved group sizes of either 30, 24, or 18 subjects, providing us with either 10, 8, or 6 subjects playing the role of each of the three player types. Subjects were recruited from the undergraduate population of the University of Pittsburgh. No subject had any prior experience with the game; subjects were only allowed to participate in a single experimental session.

In each session, subjects were seated at networked computer terminals and were isolated from one another. They received written instructions concerning the game at the beginning of each session. A copy of the instructions used in a representative session is provided in the Appendix. The written instructions were read aloud and any questions were answered. A single practice game was played to familiarize subjects with the computer interface. Subjects then began to interact with their computer terminal, providing responses when prompted. The computer program that we developed for this experiment performs the random matching of subjects, reports back to subjects on information relevant to their trading decisions, and keeps track of information such as trading decisions, goods in storage, and points earned by each subject.

We report results from 25 experimental sessions involving 636 subjects. A summary of some of the different characteristics of these experimental sessions is provided in Table 2. The primary distinctions across sessions, are the model version, A or B, and the utility value of consumption, u. These differences are indicated in Table 2. In just one session, number 16, we changed the value of the discount factor, β, from 0.90 to 0.99. The storage costs of the three goods (as reported in Table 1) remained constant across all 25 sessions. Given the parameter choices for each session, Table 2 indicates whether the pure strategy equilibrium is unique and either fundamental (F) or speculative (S), or whether both pure strategy equilibria coexist with one another (F/S). The number of subjects in each session is also given. Certain other characteristics of these experimental

TABLE 2—CHARACTERISTICS OF EXPERIMENTAL SESSIONS

Session number	Model	u	β	Equilibrium type[a]	Number of subjects	Number of games	Number of rounds[b]
1	A	20	0.90	F	30	10	66
2	A	100	0.90	S	30	10	103
3	A	100	0.90	S	24	12	111
4	A	20	0.90	F	24	12	119
5	B	20	0.90	F	24	10	123
6	B	500	0.90	F/S	24	9	83
7	B	500	0.90	F/S	24	8	88
8	B	20	0.90	F	24	12	116
9	A	100	0.90	S	30	12	142
10	A	20	0.90	F	30	13	95
11	A	100	0.90	S	30	9	93
12	A	20	0.90	F	24	9	92
13	A	100	0.90	S	24	12	89
14	A	20	0.90	F	30	21	137
15	A	15	0.90	F	30	9	103
16	A	100	0.99	S	30	1	58
17	A	100	0.90	S	24	14	93
18	A	100	0.90	S	24	8	95
19	A	100	0.90	S	24	10	136
20	A	100	0.90	S	24	13	112
21	A	100	0.90	S	24	13	98
22	A	100	0.90	S	18	8	125
23	A	100	0.90	S	18	14	92
24	A	100	0.90	S	18	14	119
25	A	100	0.90	S	30	13	107

[a] F = Fundamental, S = Speculative.
[b] Total number of trading rounds from all games played.

sessions, which concern the information subjects had available and the method used to initialize the distribution of goods in storage over player types, will be discussed in further detail later in the paper. Each of these 25 experimental sessions consisted of one or more *games* and each game consisted of a number of bilateral *trading rounds.* The number of games and trading rounds for each session is reported in Table 2. Each trading round can be described as follows.

At the beginning of each trading round, subjects are randomly paired by the computer program. Once paired, each subject's computer screen displays some information about the other player with whom the subject is matched—see Figure 1, trading round screen number 1. Subjects are asked whether they want to trade the good they have in storage for the good the player with whom they are matched has in storage. Trade occurs only if mutually agreed upon; all trades consist of one-for-one swaps of the goods the pair of players has in storage. Before making their decisions, subjects could consider some additional information that was provided on trading round screen number 1 of Figure 1. In the middle of this screen is a bar chart indicating the cumulative probability that the current game will end 1 to 10 rounds from the current round. The probabilities in this bar chart reflect the constant, one-in-ten chance (0.10 probability) that the game *will randomly end* from one round to the next. Our decision to have a random stopping rule for the game was due to two considerations. First, the Kiyotaki-Wright model envisions that players are infinitely lived. Second, the model also specifies that players discount the future using a common discount factor $\beta \in (0, 1)$. The random stopping rule is our way of dealing with both of these features of the Kiyotaki-Wright environment. As mentioned earlier, we "set" β equal to 0.90 in all but one session. To implement discounting of future payoffs by 90 percent, we chose to end each game with

Trading Round Screen # 1

You are a type 2 player. This is Round 3. Your Point Total: 110.

You earn 20 points per unit of good 2 you obtain.
It costs 1 point per round for storing good 1.
It costs 4 points per round for storing good 2.
It costs 9 points per round for storing good 3.

You currently have good 1 in storage.
You are matched with a player of type 1.
This player has good 2 in storage.
Do you want to trade your good 1 for the other player's good 2? Y/N:

Probability the game will end:

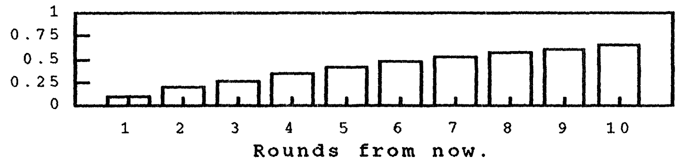

Percent of Each Type of Player Storing Each Type of Good
Historical Average as of Round 2

Type	Good 1	Good 2	Good 3
1	0.00	0.83	0.17
2	0.33	0.00	0.67
3	0.67	0.33	0.00

Trading Round Screen # 2

You are a type 2 player. This is Round 3. Your Point Total: 121.

You earn 20 points per unit of good 2 you obtain.
It costs 1 point per round for storing good 1.
It costs 4 points per round for storing good 2.
It costs 9 points per round for storing good 3.

In the last round you had good 1 in storage.
You were matched with a player of type 1.
This player had good 2 in storage.
You proposed to trade your good 1 for the other player's good 2.
The other player agreed to trade.

You received a net payoff of 11 points for the round.
Your new point total is 121.
You now have good 3 in storage.

Percent of Each Type of Player Storing Each Type of Good
Historical Average as of Round 3

Type	Good 1	Good 2	Good 3
1	0.00	0.87	0.13
2	0.25	0.00	0.75
3	0.63	0.37	0.00

The game will continue. Please wait for the next round...

FIGURE 1. ILLUSTRATION OF THE TRADING SCREENS USED IN THE EXPERIMENTS

probability $1 - \beta$, i.e., with probability 10 percent at the end of any round that has been reached.[6] The constant 0.90 probability of the game continuing into the next round implements both discounting and the stationarity associated with an infinitely lived population.

We informed subjects of this constant 10-percent probability that the game would end from one round to the next and called their attention to the bar chart revealing the cumulative probability that the game would end in future rounds.[7] The random end to the game reinforces the notion that the *time* it takes subjects to get the good they desire should also be an important consideration in their trading strategy.

At the bottom of trading round screen number 1 is a table indicating the percentage of each player type holding each good as of the end of the last trading round. In all but two sessions which will be discussed later, these percentages are *historical average* percentages based on the entire history of the current game and are updated at the end of every trading round. This information on the historical average distribution of goods by player type (the p_{ij} terms of the theory) would be useful to subjects who were attempting to calculate the relative likelihood of meeting another player who both had the good the player desired and who was willing to trade this good for the good the player currently had in storage. This type of information is assumed to be common knowledge in the theoretical Kiyotaki-Wright environment. Revealing this information to subjects was our way of attempting to implement the common knowledge assumption.

As mentioned in the introduction, our implementation of discounting and the information we provide to subjects concerning the distribution of goods held by each player type distinguishes our experimental analysis of the Kiyotaki-Wright model from the previous studies by Marimon et al. (1990) and Brown (1996). These features of our design are more conducive to speculative behavior (in environments where speculation makes sense) in that they encourage subjects to consider not only the storage cost of goods, but also to consider the time it will take them to obtain the good they desire, and the relative likelihood of meeting the player who has this good and is willing to engage in trade. Thus we view our experimental design as being somewhat closer in spirit to the Kiyotaki-Wright environment, and as providing a "best-case scenario" for speculation, in those environments designed to encourage speculative behavior.

Once all trading decisions have been made, the results of those decisions are revealed to all players on trading round screen number 2 of Figure 1. Notice that at the bottom of this screen, the information on the percentage of each type of agent holding each type of good has been updated to take account of what has occurred in the just-completed round of play.

Players carry the goods they have in storage over into the next round, *if there is a next round.* If the computer program has determined (with probability 0.10) that the just-completed round will be the last round of the game, then players see a message at the bottom of trading round screen number 2 indicating that the game has just ended. Otherwise, they see a message informing them that the game will continue with another round. If the game continues, then players are once again randomly paired with one another and are asked to make another trading decision. If the game has ended, each players' point total from that game is stored and a new game begins. Players begin every new game, including the first, with 100 points.

In most of our experimental sessions, players begin each new game with one unit of their *production good* in storage. This initialization scheme was chosen because it was

[6] Our decision to set $\beta = 0.90$ is not all that important to our predictions for which type of behavior we would expect to observe in equilibrium. In particular, if subjects did not discount future earnings at all (i.e., if $\beta = 1.0$), the various parameterizations we have chosen would remain consistent with our predictions for a fundamental (F) or speculative (S) equilibrium as indicated in Table 2.

[7] Following each trading round, the computer program drew a random number from a uniform distribution over [0, 1]. If this random draw was less than or equal to 0.10, the trading round just completed would be the last trading round of the game. Subjects were informed of the outcome of this draw only after they had made their current trading round decisions. To avoid overly long games, we also chose to end any game that exceeded 40 rounds. Subjects were not informed of this upper bound. There were only three games in all of our experimental sessions that ever reached this upper bound.

easy to explain to subjects and because it differs from any pure strategy steady-state distribution of goods over player types. Indeed, with this initialization scheme, there is always an absence of a double coincidence of wants in the first round of any game which should serve to encourage subjects to begin trading. However, we have also considered a different initialization scheme in five sessions (numbers 19–23) of model A where $u = 100$. In this alternative initialization scheme, the distribution of goods over player types at the start of each new game was made as close as possible (with a finite population) to the unique, pure strategy "speculative" equilibrium distribution of goods in storage by player type. These two different initialization schemes allow for a kind of global versus local stability analysis of the speculative equilibrium of model A, which we consider later in the paper.

At the beginning of each new game, players may have a good in storage that is different from the good they held in storage at the end of the previous game. Moreover, at the beginning of every new game, the table indicating the distribution of goods held by various types is reinitialized to reflect the initial distribution of goods by type in the new game. Thus, in each new game, players effectively "start over" with respect to the information they have available on the historical or current distribution of goods by type, and with respect to the total number of points they have earned. However, the parameterization of the game itself does not change, nor does the player's type, so that each player can draw upon his or her own past experience in playing the new game.

The number of new games that were played was determined by us and depended simply upon the time available. Subjects were recruited for two-hour sessions, but were not told in advance how many games would be played. Because our games ended randomly, the total number of games and trading rounds played varied somewhat across sessions as can be seen in Table 2; we attempted to obtain approximately 100 trading rounds per session (the average number of trading rounds per session was 104).

Once it was determined that the last game had been played, the end to the experimental session was announced. The computer program then picked one game at random from all the games played. Subjects' point totals from this game were converted into a probability of winning a \$10 prize that was in addition to the \$10 they earned for participating in the session. Each additional point subjects earned above the initial 100 points they were given increased their probability of winning the additional \$10 prize by the same amount. Subjects were informed of this mechanism for determining whether they would win the additional \$10 prize. They were further instructed that they were not competing with one another—that their probability of winning the additional \$10 prize depended only on how many additional points they were able to obtain relative to the *maximum* number of points a player of their type could have been expected to earn in the game chosen.[8] Additional \$10 prizes were then awarded based on a random choice for the cutoff probability. Our use of a binary lottery to determine actual cash payments is intended to control for subjects' differing attitudes towards risk.[9]

IV. Aggregate Experimental Results

We begin our review of the results from our experimental sessions by focusing on *aggregate* behavior by all players of a given type. Table 3 reports the frequency with which players of each type offer to trade whatever good they have in storage for the good corresponding to their type—their consumption good—over each half of all 25 sessions. Note that these are *not* the frequencies with which the three player types actually received their consumption good in trade; rather these are the frequencies with which each type *offered* to trade for their consumption good when they met another player who held this good in

[8] The maximum expected number of points takes into account a player's type, the number of trading rounds played, and the probability that a player meets his consumption good in trade.

[9] For further details concerning the use of a binary lottery procedure to control for subjects' risk preferences see, e.g., Alvin E. Roth and Michael W. K. Malouf (1979).

TABLE 3—THE FREQUENCY WITH WHICH TYPE *i* OFFERS TO TRADE FOR GOOD *i* OVER EACH HALF OF ALL SESSIONS

Session number	Type 1 offers to trade for good 1		Type 2 offers to trade for good 2		Type 3 offers to trade for good 3	
	First half	Second half	First half	Second half	First half	Second half
1	0.98	1.00	0.98	1.00	0.98	1.00
2	0.97	1.00	0.94	0.95	0.98	0.97
3	0.98	0.98	0.99	0.99	0.88	0.94
4	1.00	0.99	0.99	1.00	0.99	1.00
5	0.98	0.98	0.94	0.95	0.97	1.00
6	0.88	0.95	0.90	0.99	0.93	1.00
7	0.98	0.99	1.00	1.00	0.96	0.99
8	0.97	1.00	1.00	1.00	0.99	1.00
9	0.99	0.98	0.95	0.99	0.99	1.00
10	0.98	1.00	0.99	1.00	0.97	1.00
11	0.99	1.00	0.99	1.00	0.89	1.00
12	0.99	0.99	0.90	0.94	0.90	0.95
13	0.95	0.95	0.79	0.92	0.98	1.00
14	0.95	0.99	0.98	1.00	0.98	1.00
15	0.99	1.00	0.96	0.99	0.95	1.00
16	0.99	0.99	1.00	1.00	0.94	0.97
17	0.96	1.00	0.98	1.00	0.95	1.00
18	0.99	0.99	0.98	1.00	0.97	1.00
19	0.98	1.00	0.98	0.98	0.90	0.93
20	0.99	1.00	0.98	0.99	1.00	1.00
21	0.96	0.98	0.99	1.00	1.00	1.00
22	0.97	1.00	0.94	0.98	1.00	1.00
23	0.99	1.00	1.00	0.97	0.78	0.90
24	1.00	1.00	0.96	0.96	1.00	1.00
25	0.96	1.00	0.99	1.00	0.99	1.00
All sessions	0.98	0.99	0.97	0.98	0.96	0.99

storage. These offer frequencies indicate how well players understand their assigned roles and provide some evidence on the saliency of the monetary incentives that we provide. In theory, these offer frequencies should all be 100 percent since players are told that they receive a positive net payoff (in points) only when they successfully trade for their consumption good. Indeed we see that with few exceptions, these offer frequencies are typically close to or even equal to 1.00, especially over the second half of each session.

A. *Model A: Aggregate Results*

We now turn to the aggregate results from our experimental sessions involving model A. This version of the model is one in which the pure strategy Nash equilibrium (when it exists) is always unique, and therefore the steady-state predictions are unambiguous. We first compare aggregate results for five sessions, numbers 1, 4, 10, 12, and 14, in which we set $u = 20$ with aggregate results for five other sessions, numbers 2, 3, 9, 11, and 13, in which we set $u = 100$. The only experimental treatment condition that varies between these two groups of five sessions is the utility value of consumption, u.

Table 4 reports the frequency with which each type of player in these 10 sessions offers to trade his model A production good for the good he neither consumes nor produces. When $u = 20$, the unique pure strategy equilibrium prediction is that all three types play fundamental strategies. Type 1 players should never offer to trade their production good 2 for the higher storage cost good 3, type 2 players should *always* offer to trade their production good 3 for the lower storage cost good 1, and type 3 players should *never* offer to trade their production good 1 for the higher storage cost good 2. The offer frequencies reported in Table 4 indicate that when $u = 20$, player types 2 and 3 behave roughly in accordance with these equilibrium predictions. However, nearly one-third of type 1 players play the *speculative strategy,* offering to

TABLE 4—THE FREQUENCY WITH WHICH EACH TYPE OFFERS HIS PRODUCTION GOOD FOR THE GOOD HE NEITHER CONSUMES NOR PRODUCES OVER EACH HALF OF TEN MODEL A SESSIONS

Model A		Type 1 offers 2 for 3		Type 2 offers 3 for 1		Type 3 offers 1 for 2	
session	u =	First half	Second half	First half	Second half	First half	Second half
1	20	0.33	0.24	0.98	1.00	0.14	0.04
4	20	0.27	0.42	0.98	1.00	0.18	0.06
10	20	0.29	0.32	0.98	0.99	0.20	0.09
12	20	0.26	0.41	0.89	0.89	0.22	0.16
14	20	0.36	0.23	0.99	1.00	0.17	0.06
All	20	0.31	0.30	0.95	0.98	0.14	0.07
2	100	0.24	0.25	0.90	0.96	0.43	0.55
3	100	0.21	0.38	0.97	0.99	0.11	0.09
9	100	0.22	0.19	0.93	0.96	0.18	0.14
11	100	0.50	0.58	0.93	1.00	0.24	0.30
13	100	0.66	0.68	0.83	0.84	0.23	0.21
All	100	0.36	0.36	0.90	0.95	0.24	0.25

trade their production good 2 for the relatively higher storage cost good 3.

When $u = 100$, the equilibrium predictions for player types 2 and 3 are the same as when $u = 20$, but the equilibrium prediction for type 1 players is that these players will *always* speculate by offering to trade their production good 2 for the higher storage cost good 3. In Table 4 we see that in the five sessions where $u = 100$, the aggregate frequencies with which type 1s offer to trade good 2 for good 3 are not much different from those reported for the sessions where $u = 20$. While there is some variance in these offer frequencies across sessions, we note that the weighted average frequency with which type 1 players offer to trade good 2 for good 3—36 percent—does not change from the first to the second half of all sessions where $u = 100$. Nonparametric, robust rank-order tests were conducted using the five offer frequencies for player type 1s in the two treatments where $u = 20$ and $u = 100$.[10] The null hypothesis is that there is no difference in the distribution of type 1 offer frequencies between the two treatments. We find that we *cannot reject* this null hypothesis (at the 0.05 significance level) for either the first or the second half of all sessions.[11] We conclude that, at the aggregate level, player type 1s do not respond to variations in the utility value of consumption (u) that were intended to elicit fundamental or speculative behavior.

Using robust rank-order tests on the offer frequencies of player type 2s over each half of a session, we are also unable to reject (at the 0.05 level) the null hypothesis of no difference in these offer frequencies between the two treatments. However, as noted above, the behavior of player types 2 and 3 should not change as we vary the value of u. When we apply the robust rank-order test to the offer frequencies of player type 3s over each half of a session, we find that for the second half of all model A sessions, we *can reject* (at the 0.05 level) the null hypothesis of no difference in offer frequencies between the two treatments in favor of the alternative that offer frequencies are *higher* in sessions where $u = 100$ than in sessions where $u = 20$. This finding runs counter to the theoretical prediction that type 3 behavior will be unaffected by variations in the value of u.

[10] See Sidney Siegel and N. John Castellan, Jr. (1988 pp. 137–44) for an explanation of this nonparametric test.

[11] If we compared the three *lowest* type 1 offer frequencies from the second half of sessions where $u = 20$—sessions 1, 10, and 14—with the three *highest* type 1 offer frequencies from the sessions where $u = 100$—sessions 3, 11, and 13—we would be able to reject the null hypothesis. Hence the finding reported above depends on the observed variance across sessions in type 1 offer frequencies under both treatments.

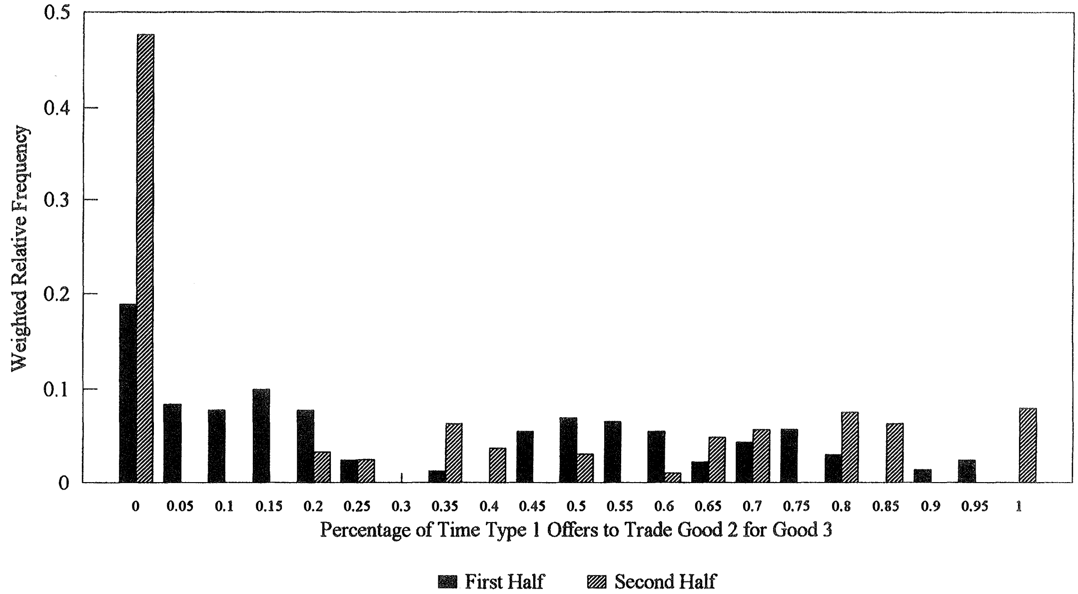

FIGURE 2A. HISTOGRAM OF THE PERCENTAGE OF SPECULATIVE PLAY BY TYPE 1 PLAYERS EACH HALF OF SESSIONS 1, 4, 10, 12, AND 14 WHERE $u = 20$, 46 TYPE 1 SUBJECTS

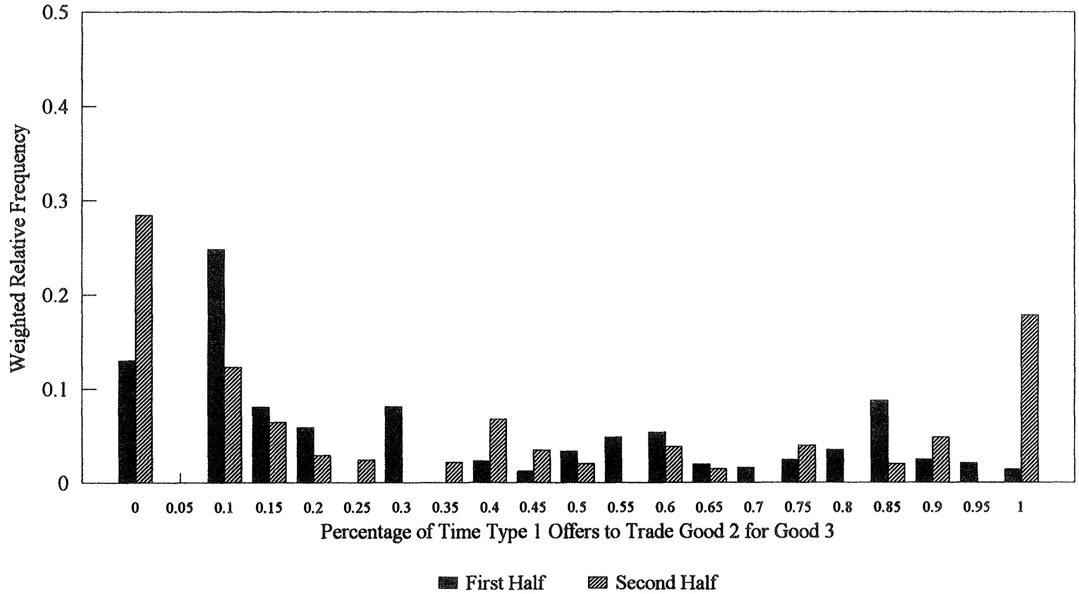

FIGURE 2B. HISTOGRAM OF THE PERCENTAGE OF SPECULATIVE PLAY BY TYPE 1 PLAYERS EACH HALF OF SESSIONS 2, 3, 9, 11, AND 13 WHERE $u = 100$, 46 TYPE 1 SUBJECTS

Examination of the aggregate frequencies reported in Table 4 leads naturally to the question of whether these frequencies reflect the distribution of individual strategies. Figure 2A presents a weighted relative frequency histogram of the percentage of speculative play by type 1 players over each half of the five sessions where $u = 20$ and Figure 2B does the same for the type 1 players in the five sessions where $u = 100$. The percentage of speculative type 1 offers of good 2 for good 3 range from 0, indicating the play of the fundamental pure strategy on up to 1, indicating the play of the speculative pure strategy. In both treatments, the percentage of type 1 players playing either fundamental or speculative pure strategies grows from the first to the second half of the experimental sessions, with the percentage of those playing the fundamental pure strategy (or something close to it) representing the most frequently observed outcome. Similar results, not illustrated here, are found for type 3 players in both treatments. By contrast, nearly all type 2 players always played the fundamental pure strategy of offering good 3 for good 1. We conclude that there may indeed be considerable heterogeneity among players assigned to play the role of certain types; by the second half of the sessions of either treatment, more than 50 percent of type 1 players play *pure* strategies at least 90 percent of the time, while the remainder play some *mixture* of pure strategies, though there appears to be no consensus on the choice of the mixed strategy probability.

Another concern that arises with respect to the results reported in Table 4 is whether the conditions that guaranteed the existence of a fundamental or speculative pure strategy equilibrium were, in fact, in place during the course of the experimental sessions involving model A. Recall that in the unique pure strategy Nash equilibrium of model A, we expect to observe *fundamental* behavior by type 1 players whenever the inequalities in (1) are satisfied. Thus, it is important to check whether in treatments where $u = 20$, it was the case that $(p_{31} - p_{21}) < 5/6$ and in treatments where $u = 100$ (the speculative parameterization where type 1s are called on trade good 2 for good 3 in equilibrium), it was the case that $p_{31} - p_{21} > 1/6$. We found that the average value of the difference $p_{31} - p_{21}$ in all sessions involving model A was always *less* than $5/6$ and was always *greater* than $1/6$. Thus, the theoretical conditions were in place so that on average, a type 1 player's steady-state best response would have been to play the fundamental strategy when $u = 20$ and would have been to play the speculative strategy when $u = 100$. Whether the type 1 players *recognized* the different incentives that were present in the two different environments will be addressed later in Section V when we consider individual learning behavior.

When $u = 100$, the incentives for all players to play the predicted equilibrium strategies as opposed to all players playing fundamental strategies, are quite large. If all players play fundamental strategies, the expected pattern of trade is such that the steady-state probability that each player type gets his consumption good in trade in any round is $1/6$. However, if type 1s speculate while type 2s and 3s play fundamental then the steady-state probability that each player type gets his consumption good in trade in any round is $(1/6)\sqrt{2}$, an increase of more than 40 percent.

Do subjects respond to these incentives in making strategic choices? Figure 3 provides a comparison of the strategy choices of the three player types in the 10 model A sessions discussed in Table 4 where $\beta = 0.90$ and $u = 20$ or 100. Our comparison is based on an "efficiency score" that we constructed for each player type in every round of these sessions. This efficiency score measures the frequency with which players of a given type play "best responses" in the interesting trading situations where: (1) they have the opportunity to trade for a good that does not correspond to their type, and (2) the good they face in trade is different from the good they have in storage.

In making these best-response calculations, we assumed that each player played according to the steady-state predictions. That is, we imagined that type 1 players played the fundamental or the speculative strategy depending on whether the actual historical distribution of goods over types, which they observed on their computer screens, warranted the play of either strategy. We assumed that player types 2 and 3 also played according to the equilibrium, steady-state predictions by adhering to fundamental strategies regardless of the value of u or the information on the historical distribution of goods over types.[12] For each round, we

[12] We could not think of another way to model best-response behavior as subjects in these sessions did not know the distribution of *strategies* that subjects were playing.

Therefore, we chose to focus on the steady-state strategies.

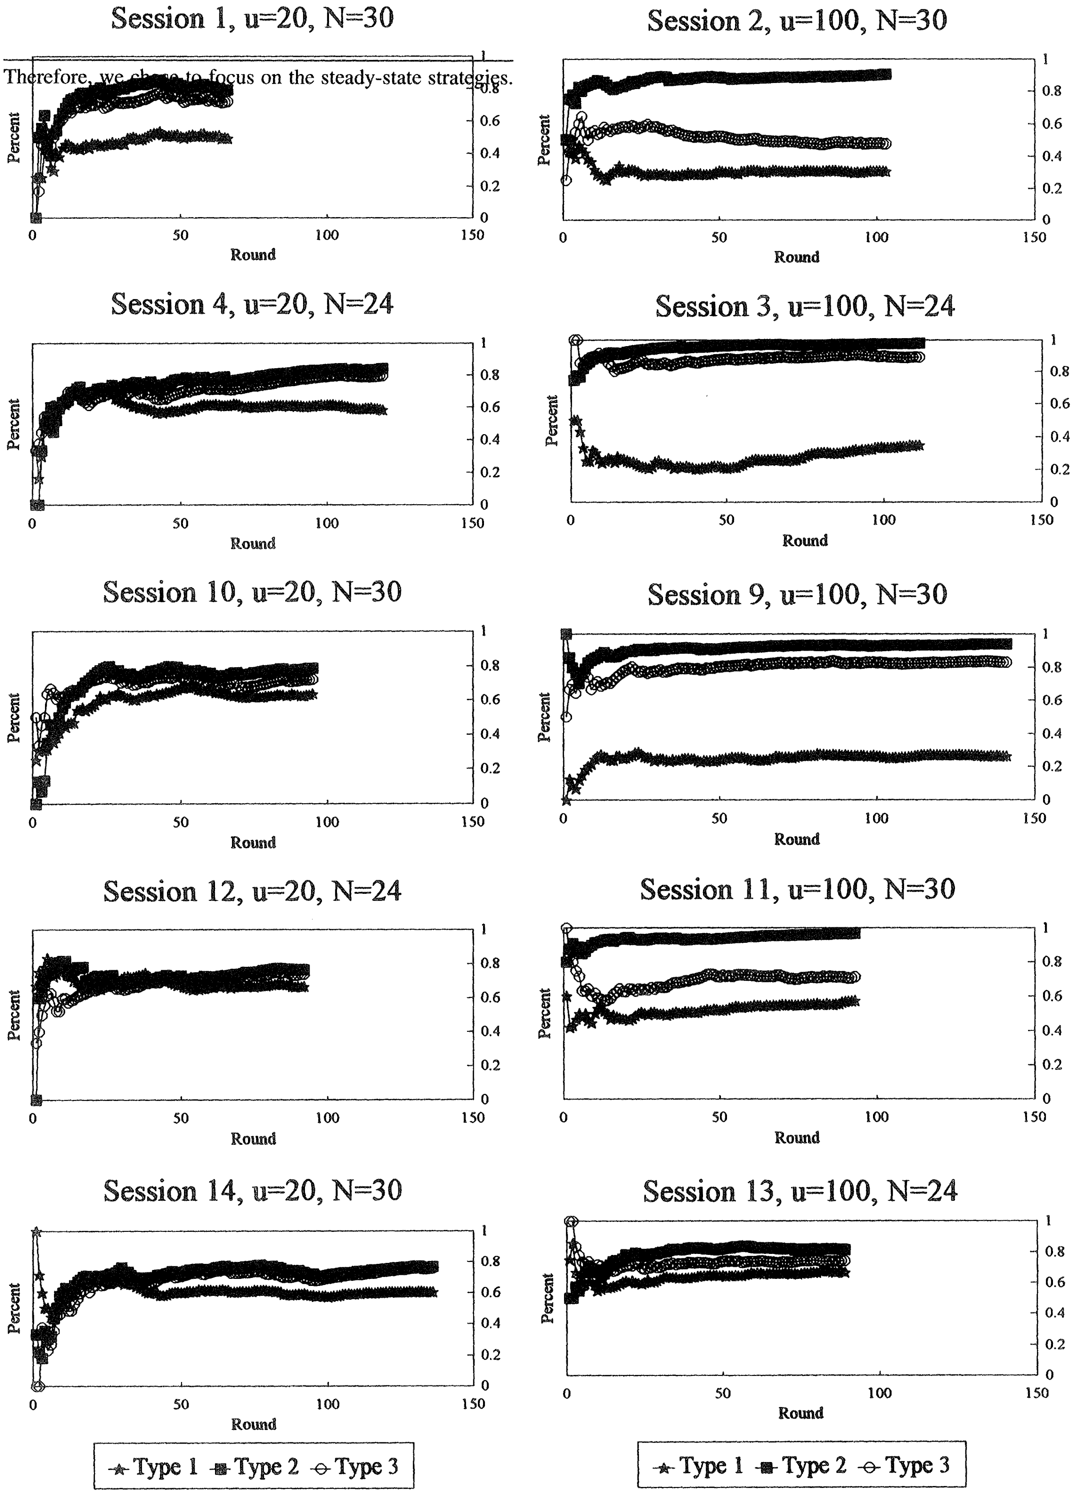

FIGURE 3. "BEST-RESPONSE" EFFICIENCY SCORES FOR MODEL A SESSIONS

calculated the average frequency with which players of a given type in the strategically interesting trading situations have played best responses using all data from the beginning of the session to the current round. These best-response efficiency scores for the model A sessions are presented in Figure 3. We see that in most of the sessions where $u = 20$, player type 1s are approximately as efficient as player types 2 and 3 at playing best responses to historical frequencies. By contrast, in most sessions where $u = 100$, player type 1s are substantially less efficient at playing best responses than are player types 2 and 3. The lower efficiency scores for player type 1s in the $u = 100$ sessions arise from the refusal of a large number of type 1 players to engage in the speculative strategy in these sessions.

B. *Robustness of the Aggregate Results for Model A*

One feature of our experimental design is that subjects in each session play more than one game. This provides subjects with direct experience of the random determination of the end of a game, a feature that introduces subjects to the discounting that must take place if there is to be a trade-off between storage costs and marketability considerations. Subjects began each new game in the model A sessions reported in Table 4 with their production good in storage. Thus, at the beginning of each new game, only type 3 players are storing good 1. This initial condition, *if it persisted,* would make speculative trades by type 1 players (trades of good 2 for good 3) profitable in model A sessions where $u = 20$. However, fundamental play by type 2 and 3 players imply that this initial condition will not persist. Nevertheless, since we reinitialized inventory holdings at the start of each new game, it is possible that our reinitialization may have induced the play of speculative strategies by type 1 players in environments ($u = 20$) where in equilibrium, they should not speculate. Additionally, we wanted to consider whether the frequency with which games end and new games begin—the expected number of trading rounds per game when $\beta = 0.90$ is 10—may have impeded type 1 players' adoption of the speculative strategy in environments ($u = 100$) where this strategy is the steady-state best response.

To check whether our reinitialization scheme may have led to these kinds of biases in our results, we conducted two additional experimental sessions, involving model A, numbers 15 and 16. In session 15, we set $u = 15$ rather than $u = 20$. When $u = 15$ in our parameterization of model A, inequality (1) becomes $p_{31} - p_{21} < 10/9$, a condition that *always holds* irrespective of our reinitialization scheme. Thus, when $u = 15$, a type 1 player's steady-state best response is to always play the fundamental strategy.

Results from session 15 are reported in Table 5. A comparison of these results with those reported in Table 4 for sessions 1, 4, 10, 12, and 14 where $u = 20$ suggests that there is little difference in behavior, especially over the second half of a session.

Session 16 represents an effort to remove any effect of reinitialization by *raising the discount factor,* β, from 0.90 to 0.99, thereby increasing the expected length of a game from 10 to 100 rounds.[13] In this session a single game was played that lasted for 58 rounds. This represents an increase of almost 50 percent in the number of rounds played in any game of any other session. In this session, u was set equal to 100; therefore, the equilibrium prediction is that type 1 players will adopt the speculative strategy. The results from this 58-round game (session 16) are reported in Table 6. A comparison of Table 6 with the results reported in Table 4 for sessions 2, 3, 9, 11, and 13 where u also equaled 100 shows that there is little difference in the behavior of all three types. In particular, player type 1s in session 16 where $\beta = 0.99$ do not appear to play speculative strategies any more frequently than they do in sessions where $\beta = 0.90$. Furthermore, a histogram depicting the weighted relative frequency with which type 1s played the two pure strategies (not shown here) in session 16 is similar to the one shown in Figure 2B. We conclude that increasing the expected number of rounds that will be played in a game does not appear to promote coordination on a Nash equilibrium.

All of the aggregate results for model A discussed thus far involve an initialization scheme in which each player begins each new game with the good that his type produces in storage. As mentioned earlier, an alternative initialization scheme is to begin each new game with a

[13] For session 16, we also increased the maximum number of rounds allowed from 40 to 200.

TABLE 5—THE FREQUENCY WITH WHICH EACH TYPE OFFERS HIS PRODUCTION GOOD FOR THE GOOD HE NEITHER CONSUMES NOR PRODUCES OVER EACH HALF OF A MODEL A SESSION WHERE $u = 15$

Model A	Type 1 offers 2 for 3		Type 2 offers 3 for 1		Type 3 offers 1 for 2	
$u =$	First half	Second half	First half	Second half	First half	Second half
15	0.19	0.32	0.90	0.91	0.19	0.06

TABLE 6—THE FREQUENCY WITH WHICH EACH TYPE OFFERS HIS PRODUCTION GOOD FOR THE GOOD HE NEITHER CONSUMES NOR PRODUCES OVER EACH HALF OF A MODEL A SESSION WHERE $u = 100$ AND $\beta = 0.99$

Model A	Type 1 offers 2 for 3		Type 2 offers 3 for 1		Type 3 offers 1 for 2	
$u =$	First half	Second half	First half	Second half	First half	Second half
100	0.16	0.20	1.0	1.0	0.20	0.14

distribution of goods over player types that is as close as possible (with a finite population) to the *steady-state* distribution. We executed this alternative initialization scheme for the model A treatment where $u = 100$, as we thought this initialization scheme might help to promote speculation by player type 1s and might also reduce the frequency with which player type 3s engage in speculative trades. We conducted five model A sessions, numbers 19–23, where $u = 100$ and the initial distribution of goods over types was made as close as possible to the speculative steady-state equilibrium distribution. In these sessions, each player type began each new game with either his production good or the good he neither produces nor consumes in storage so that in the aggregate, the distribution of goods over types was close to the speculative steady-state distribution.[14]

Table 7 presents the aggregate frequency with which each type offered to trade his production good for the good he neither consumes nor produces in these five sessions. A comparison of the five offer frequencies reported in Table 7 with those reported in Table 4 for sessions where $u = 100$ and players began each game with their production good in storage reveals no significant differences. While there is, again, some variance in the offer frequencies across sessions reported in Table 7, we see that, on average, type 1 players do not speculate any more frequently, nor do type 3 players speculate much less frequently, when the aggregate distribution of goods across player types is initially very close to the speculative steady-state distribution (as compared with when it is not, as in Table 4).

Robust rank-order tests confirm these findings. If we compare the five aggregate offer frequencies for each player type over each half of a session as reported in Table 7 with the corresponding set of five offer frequencies reported for the five sessions of Table 4 where $u = 100$, we can never reject (at the 0.05 significance level), the null hypothesis that both sets of aggregate offer frequencies are drawn from the same distribution. Furthermore, the distribution of speculative offers by types 1 and 3 in the sessions reported in Table 7 is similar to that illustrated in Figure 2B, with the largest fraction of type 1 and 3 players adhering to a fundamental pure strategy by the second half of all five sessions. We conclude that subject behavior is invariant to the two different initialization schemes we have considered.

As a final robustness check on the aggregate results we obtained for model A, we altered the information that we reported to subjects concerning the distribution of goods held in storage by each player type. In all of the experimental sessions discussed so far, subjects received information on the historical average distribution of goods held by each player type over all rounds of the current game as of the last trading round. This information was intended to serve as a proxy for

[14] The initial proportions we used depended on the number of subjects available. In sessions 19–21, where we had 24 subjects (8 of each type), we set $p_{12} = 6/8$, $p_{23} = 3/8$, and $p_{31} = 8/8$. In sessions 22–23 where we had 18 subjects (6 of each type), we set $p_{12} = 4/6$, $p_{23} = 2/6$, and $p_{31} = 6/6$. In model A, the unique speculative steady-state distribution is: $p_{12} = 0.71$, $p_{23} = 0.41$, and $p_{31} = 1.0$.

TABLE 7—THE FREQUENCY WITH WHICH EACH TYPE OFFERS HIS PRODUCTION GOOD FOR THE GOOD HE NEITHER CONSUMES NOR PRODUCES OVER EACH HALF OF MODEL A SESSIONS WHERE THE INITIAL DISTRIBUTION OF GOODS IS CLOSE TO THE SPECULATIVE EQUILIBRIUM DISTRIBUTION

Model A		Type 1 offers 2 for 3		Type 2 offers 3 for 1		Type 3 offers 1 for 2	
session	$u =$	First half	Second half	First half	Second half	First half	Second half
19	100	0.13	0.18	0.98	0.97	0.29	0.29
20	100	0.38	0.65	0.95	0.95	0.17	0.14
21	100	0.48	0.57	0.96	1.00	0.13	0.14
22	100	0.08	0.24	0.92	0.98	0.12	0.02
23	100	0.06	0.32	0.93	0.97	0.25	0.18
All	100	0.23	0.37	0.95	0.96	0.20	0.16

information on strategic behavior by each player type. However, subjects may have had difficulty using this information to infer the distribution of strategies that were being played. Therefore, we conducted two sessions of model A with $u = 100$, numbers 17 and 18, in which we gave subjects additional information. In these sessions, we reported the historical average frequency with which type i players, with good $j \neq i$ in storage, offered to trade good j for good $k \neq j$ in the current game, as of the last trading round. These 12 "strategy frequencies" were presented on the second trading round screen. This strategy information was *in addition* to the information given on the historical average frequency with which each player type was storing each good, thus allowing us to compare and contrast the effect of providing the additional information on strategic behavior. The aggregate offer frequencies from sessions 17 and 18, where subjects had information on strategic behavior, are given in Table 8.

Comparing the offer frequencies in Table 8 with the offer frequencies for the five model A sessions reported in Table 4 where $u = 100$ but no strategy information was given, we do not see much difference. In particular, we continue to find that less than half of type 1 players play the speculative strategy of trading good 2 for good 3 and a roughly constant number of type 3 players choose to play the speculative strategy of trading good 1 for good 2.

In two other model A sessions, numbers 24 and 25, we returned to our standard experimental design, where we do not reveal strategy information, and we considered a different modification to the information we provided to subjects. Rather than providing information on the *historical average* distribution of goods held in storage by player type, we instead reported the *current* round-to-round distribution of goods in storage by player type, and we updated this information at the end of every round. In the later rounds of a game, this information on the current distribution of goods by type might be a better predictor of the goods a player is likely to encounter in each new trading round, and may therefore facilitate the play of speculative strategies. The aggregate offer frequencies from sessions 24–25, where players were given the current distribution of goods by type, are given in Table 9.

Again, we see that there is not much difference between the aggregate offer frequencies in this new treatment as compared with the offer frequencies of our standard treatment, as reported in Table 4, where subjects had information on the historical average distribution of goods held by each player type. We also note that in the two sessions 24–25 reported in Table 9, as well as in the two sessions 17–18 reported in Table 8, the distribution of strategies across player types was again somewhat heterogeneous, with the largest fraction of players of each type adhering to fundamental pure strategies by the second half of all sessions.

The findings of this section suggest that the aggregate results we obtained using our "standard" experimental design for model A, as reported in Table 4, are robust to several different modifications. In particular, the results are not an artifact of having many games in which subjects' inventory holdings get reinitialized, nor are they affected by initializing the distribution of goods close to the steady-state distribution, nor by the addition of information on strategic behavior or the current distribution of goods held in storage. We will therefore use our standard experimental design, in which $\beta = 0.9$, players begin each new game with their produc-

TABLE 8—THE FREQUENCY WITH WHICH EACH TYPE OFFERS HIS PRODUCTION GOOD FOR THE GOOD HE NEITHER CONSUMES NOR PRODUCES OVER EACH HALF OF TWO MODEL A SESSIONS IN WHICH STRATEGY INFORMATION WAS GIVEN

Model A		Type 1 offers 2 for 3		Type 2 offers 3 for 1		Type 3 offers 1 for 2	
session	$u =$	First half	Second half	First half	Second half	First half	Second half
17	100	0.27	0.36	0.90	0.99	0.14	0.16
18	100	0.49	0.49	0.98	1.00	0.14	0.14
All	100	0.36	0.42	0.94	0.99	0.14	0.15

TABLE 9—THE FREQUENCY WITH WHICH EACH TYPE OFFERS HIS PRODUCTION GOOD FOR THE GOOD HE NEITHER CONSUMES NOR PRODUCES OVER EACH HALF OF TWO MODEL A SESSIONS IN WHICH THE CURRENT DISTRIBUTION OF GOODS WAS PROVIDED

Model A		Type 1 offers 2 for 3		Type 2 offers 3 for 1		Type 3 offers 1 for 2	
session	$u =$	First half	Second half	First half	Second half	First half	Second half
24	100	0.26	0.33	0.91	0.90	0.21	0.18
25	100	0.48	0.52	1.00	0.99	0.19	0.09
All	100	0.39	0.44	0.97	0.96	0.20	0.13

tion good in storage and are given information on the historical average distribution of goods over player types, to explore some of the equilibrium predictions that emerge from model B.

C. *Model B: Aggregate Results*

In contrast to model A, in model B there *always* exists, (for all valid parameterizations of the model), a pure strategy Nash equilibrium where all types play fundamental strategies. For a subset of this same parameter space, there also exists a pure strategy Nash equilibrium where player types 2 and 3 play speculative strategies, while type 1 players continue to play fundamental. We report here aggregate results from two sessions of model B, sessions 5 and 8, where we set $u = 20$ and we compare these results with results we obtained from two other sessions of model B, sessions 6 and 7, where we set $u = 500$. The latter two sessions are the only ones in which speculative behavior on the part of player types 2 and 3 is potentially consistent with the theory. Table 10 reports the frequency with which each type of player offers to trade his model B production good for the good that he neither consumes nor produces.

The frequencies reported in Table 10 indicate that in sessions 5 and 8 where $u = 20$, the majority of all three player types follow fundamental trading strategies in accordance with the theoretical equilibrium prediction. Moreover, this preference for the fundamental strategy appears to become stronger in the second half of each of these two sessions. In particular, we see that most type 1 players offer to trade their production good 3 for the lower storage cost good 1, most type 2 players refuse to trade their production good 1 for the higher storage cost good 3 and most type 3 players offer to trade their production good 2 for the lower storage cost good 1. Finally, we note that the tendency of all types to play fundamental strategies in sessions 5 and 8 of model B appears to be somewhat stronger (especially in the second half of each session) than was the case for the model A sessions reported in Table 4 where $u = 20$ and the fundamental equilibrium was also the unique Nash equilibrium.

In the two sessions 6 and 7 where $u = 500$, we see an increased tendency on the part of *all three* player types to play speculative strategies relative to the case where $u = 20$. However, the frequency of speculative play by player types 2 and 3 remains far from the 100-percent frequency that is required for a speculative Nash equilibrium. While the behavior of player types 2 and 3 is in the direction of the speculative equilibrium, the behavior of type 1 players is not; for type 1 players, the fundamental strategy is always the steady-state best response, regardless of the parameterization of model B.[15] Thus, when $u = 500$, there appears to be no evidence that players move

[15] Note the similarity between the response of player type 1s in model B to an increase in u and the response of player type 3s in model A to an increase in u.

TABLE 10—THE FREQUENCY WITH WHICH EACH TYPE OFFERS HIS PRODUCTION GOOD FOR THE GOOD HE NEITHER CONSUMES NOR PRODUCES OVER EACH HALF OF FOUR MODEL B SESSIONS

Model B		Type 1 offers 3 for 2		Type 2 offers 1 for 3		Type 3 offers 2 for 1	
session	u =	First half	Second half	First half	Second half	First half	Second half
5	20	0.98	0.97	0.13	0.06	0.85	0.90
8	20	0.92	0.99	0.06	0.02	0.94	1.00
All	20	0.95	0.98	0.10	0.04	0.90	0.95
6	500	0.70	0.88	0.20	0.10	0.89	0.86
7	500	0.91	0.86	0.10	0.22	0.69	0.75
All	500	0.82	0.87	0.13	0.16	0.75	0.81

closer to the speculative equilibrium of model B, nor is there any evidence that players move toward a mixed strategy equilibrium, as the behavior of type 1 players is inconsistent with either of these two conjectures.[16] As we did for model A, we also checked whether the conditions were in place to guarantee the uniqueness of the fundamental equilibrium of model B or to allow the speculative equilibrium of model B to coexist with the fundamental equilibrium. It is easily verified that the first inequality in (2) that is necessary for a speculative equilibrium to coexist with the fundamental equilibrium is *always violated* when u = 20 since the difference, $p_{32} - p_{12}$, cannot exceed 1.0 by definition.

Furthermore, we found that the second inequality in (2) involving p_{23} also fails to hold, on average, in the two sessions 5 and 8 where $u = 20$. Therefore, in both sessions 5 and 8 the conditions were such that the fundamental equilibrium could be regarded as the unique Nash equilibrium. Similarly, we verified that in the sessions where u = 500 (sessions 6 and 7), the inequalities in (3) were *satisfied,* on average, over all rounds of these two sessions. Therefore, in both sessions 6 and 7, the conditions necessary for the speculative equilibrium to coexist with the fundamental equilibrium were in place, on average, over all rounds of these two sessions.

The multiplicity of pure strategy steady-state equilibria in the model B environment, in particular, the existence of a fundamental pure strategy equilibrium for all parameter values, makes alternative experimental treatments intended to encourage the play of speculative strategies less useful in the model B environment. For this reason, we have not pursued such alternative treatments as we did for model A. We simply note that our aggregate findings for model A appear to carry over to the model B environment with its different distribution of production goods over player types.

[16] We do not have enough experimental sessions involving model B to test for aggregate, session-level differences between the two treatments (u = 20, u = 500) using robust rank-order tests.

V. Learning Behavior

In this section we ask whether players are in fact responding to *marketability* considerations, as the theory implies, or whether there may be some other explanation for players' behavior. In an effort to answer this question we focus on the strategic behavior of players who must decide whether to trade for a good that: (1) is not their consumption good, and (2) differs from the good they currently hold in storage.

Denote by s^i_{jk} the strategy chosen by player type i who has good $j \neq i$ in storage and who meets a player with good $k \neq j, i$. For simplicity, we suppress all references to time. Let $s^i_{jk} = 0$ if player type i plays the *fundamental* strategy and let $s^i_{jk} = 1$ if player type i plays the *speculative* strategy. We wish to relate the probability that s^i_{jk} = 0 (that type i with good j plays "fundamental") to two independent variables.

The first independent variable, NETPAY_{jk}, is defined by:

$$\text{NETPAY}_{jk} = (\text{success}_{ji} - \text{fail}_{ji}) - (\text{success}_{ki} - \text{fail}_{ki}).$$

Here $(\text{success}_{ji} - \text{fail}_{ji})$ measures the difference between the number of times, up to the current round of a session, a type i player with good j successfully traded good j for good i, and the number of times the type i player failed to trade good j for good i.[17] This difference measures the *rewards* the type i player experienced when holding good j. The second difference ($\text{success}_{ki} - \text{fail}_{ki}$) is calculated similarly, and measures the rewards the type i player experienced when holding good $k \neq j$. The NETPAY_{jk} variable captures the past success of the strategy of storing good j relative to the past success of the strategy of storing good k. This variable serves as a proxy for *reinforcement-based* effects on player behavior. Such reinforcement effects are the primary mechanism through which individual learning is modeled in a variety of adaptive learning schemes, including the classifier system of Marimon et al. (1990). Indeed, information about past payoff performance is typically the *only* kind of information that agents have available in reinforcement-based learning models.

The second independent variable measures the marketability of good j relative to good k. The construction of this marketability variable makes use of the actual historical average frequencies that players observed on their computer screens in theoretically relevant ways. We constructed these marketability variables for player type 1s in model A, and for player types 2 and 3 in model B.

In model A, type 1 players are concerned with the marketability of good 2 relative to good 3. A type i player with good j knows that if type j holds good i, that player will accept good j in trade. Therefore, p_{ji} is a measure of the relative marketability of good j from type i's perspective. Furthermore, from (1) we know that for type 1 players, the difference $(p_{31} - p_{21})$ must be less than a certain threshold value—5/6 when $u = 20$ and 1/6 when $u = 100$—for fundamental behavior by player type 1s to be a best response to the fundamental behavior of player type 2s and 3s. Therefore, we constructed the following "marketability variables" for player type 1s

$$P21P3156 = p_{21} - p_{31} + \tfrac{5}{6},$$

$$P21P3116 = p_{21} - p_{31} + \tfrac{1}{6},$$

using the actual historical frequency values for p_{21} and p_{31} that were revealed to all subjects on their computer screens prior to every trading decision. By construction, player type 1s should play the *fundamental* strategy in treatments where $u = 20$ whenever $P21P3156 > 0$ and in treatments where $u = 100$ whenever $P21P3116 > 0$; they should play the speculative strategy whenever these marketability variables are negative.

Similarly, in model B, we constructed marketability variables for player types 2 and 3 in accordance with the inequalities given in (2–3). For type 2 players in model B, we constructed the marketability variables

$$(4) \qquad P12P3243 = p_{12} - p_{32} + \tfrac{4}{3},$$

$$(5) \qquad P12P32475 = p_{12} - p_{32} + \tfrac{4}{75},$$

and for type 3 players in model B we constructed the marketability variables:

$$(6) \qquad P2312 = \tfrac{1}{2} - p_{23},$$

$$(7) \qquad P23150 = \tfrac{1}{50} - p_{23}.$$

Expressions (4) and (6) are the marketability variables for types 2 and 3, respectively, in model B when $u = 20$. Expressions (5) and (7) are the corresponding variables in model B for types 2 and 3 when $u = 500$. By construction, types 2 and 3 should play *fundamental* strategies when these two marketability variables are positive, and when these two variables are negative, the speculative equilibrium *coexists* with the fundamental equilibrium.

Using these measures, we estimate logit regressions in which the probability that $s^i_{jk} = 0$ depends on: (1) the relative past success of this strategy as measured by the NETPAY_{jk} variable, (2) the appropriate marketability

[17] For simplicity, we again suppress any reference to time. Note also our assumption that players only pay attention to the good they had in storage when they successfully traded for the good they desire; we assume they do not assign any credit to the chain of goods they may have held in storage prior to holding good j.

TABLE 11—REGRESSION RESULTS FOR PLAYER TYPE 1S IN MODEL A

Independent variable[a]	(a) $u = 20$				(b) $u = 100$				(c) $u = 100$ and initialization at steady state			
	s^1_{23}	s^1_{23}	s^1_{32}	s^1_{32}	s^1_{23}	s^1_{23}	s^1_{32}	s^1_{32}	s^1_{23}	s^1_{23}	s^1_{32}	s^1_{32}
Constant	3.2996***	2.3652***	4.7560**	4.5877**	0.0151	−0.4279	−0.7199	−0.6938	2.0842***	1.9538***	2.1560**	1.0971
	(0.5500)	(0.5094)	(2.1990)	(2.1835)	(0.3120)	(0.2993)	(0.9193)	(0.9132)	(0.3782)	(0.3360)	(0.9181)	(0.7714)
NETPAY_{23}	0.0583***	—	0.0475*	—	0.0862***	—	−0.0065	—	0.0823***	—	0.1034**	—
	(0.0104)		(0.0281)		(0.0108)		(0.0265)		(0.0181)		(0.0409)	
$P21P3156$	1.2389***	1.0731***	3.1758**	3.4913**								
	(0.3253)	(0.3166)	(1.5398)	(1.5244)								
$P21P3116$					0.3335	0.3363	−0.0018	−0.0010	1.3479	1.7116	1.1518	1.0090
					(0.3031)	(0.2962)	(1.2363)	(1.2363)	(1.2488)	(1.2349)	(2.3200)	(2.2584)
$\Pr > \chi^2$ [b]	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.9670	0.9303	0.0001	0.0001	0.0001	0.4496
Observations	992	992	134	134	1,150	1,150	183	183	752	752	186	186

Notes: Standard errors are reported in parentheses.

[a] The coefficient estimates on four session dummies have been omitted.

[b] p-values from a Wald test of the hypothesis that the coefficient estimates on the independent variables (other than the intercept) are all zero.

* Significantly different from 0 at the 10-percent level.

** Significantly different from 0 at the 5-percent level.

*** Significantly different from 0 at the 1-percent level.

variable, (3) a constant term, and (4) session dummy variables for observations from different experimental sessions within a treatment.

A. *Player Type 1s in Model A*

Panels (a) and (b) of Table 11 report logit regression results for our learning model using pooled data for type 1 players in model A. The data used in the regressions reported in Table 11(a) are from the five sessions of our "standard" experimental design where $u = 20$ (session numbers 1, 4, 10, 12, and 14) and the data used in the regressions reported in Table 11(b) are from the five sessions of our standard experimental design where $u = 100$ (session numbers 2, 3, 9, 11, and 13). The type 1 player strategies are divided up according to whether the player had good 2 in storage and faced good 3, s^1_{23}, or had good 3 in storage and faced good 2, s^1_{32}. The logit regressions in Table 11 estimate the probability with which type 1 players in either trading situation play the *fundamental strategy,* i.e., $\Pr(s^1_{ij} = 0)$; for this reason, we expect the coefficient estimates on the netpay and marketability variables to be positive due to the way in which these independent variables have been defined. For example, in a regression of s^1_{23} on NETPAY_{23} and $P21P3156$, a positive coefficient on the NETPAY_{23} variable indicates that type 1 players play the fundamental strategy of refusing to trade good 2 for good 3, $s^1_{23} = 0$, with a greater probability when $\text{NETPAY}_{23} > 0$ and with a lower probability when $\text{NETPAY}_{23} < 0$, in accordance with the reinforcement learning hypothesis. Similarly, a positive coefficient on the $P21P3156$ marketability variable suggests that type 1 players are more likely to play the fundamental strategy when $P21P3156 > 0$ and are less likely to play the fundamental strategy when $P21P3156 < 0$ in accordance with the predictions of the theory. Note also that we have considered two different specifications of our learning model: one specification includes the NETPAY_{23} variable and the other does not.

Consider first, the strategic behavior of type 1 players who are storing their production good 2 when they meet another player with good 3, i.e., s^1_{23}. We see that regardless of whether $u = 20$ [Table 11(a)] or $u = 100$ [Table 11(b)], the coefficient on the reinforcement variable, NETPAY_{23}, is always positive and significantly different from zero, indicating that type 1 players are more likely to play the fundamental strategy of refusing to trade good 2 for good 3 the higher is the net payoff from holding good 2 relative to the net payoff from holding good 3. We also see that when $u = 20$, the coefficient on the relevant marketability variable, $P21P3156$, is positive and significantly different from zero, indicating that type 1 players are more likely to play the fundamental strategy when $P21P3156 > 0$ and are less likely to play this strategy when $P21P3156 < 0$ as the theory predicts. By contrast, when $u = 100$, the coefficient on the relevant marketability variable, $P21P3116$, while positive, is not significantly different from zero, indicating that in the environment where the theory predicts that player type 1s will play the speculative strategy in equilibrium, these player 1s are, in fact, ignoring information about marketability that is critical to their adoption of the speculative strategy. Notice further that the coefficient on the marketability variable, $P21P3116$, remains insignificant when we eliminate the NETPAY_{23} variable from the model specification. These results provide some further insight into our aggregate finding that the majority of player type 1s refuse to play the speculative strategy in the environment ($u = 100$), where speculative play is a best response.

Similar results are found for the strategic behavior of type 1 players who have good 3 in storage when they meet another player with good 2, i.e., s^1_{32}. Here the type 1 player's fundamental strategy, $s^1_{32} = 0$, is to *offer to trade* good 3 for good 2, while the speculative strategy is to refuse such trades.[18] Again we see that when $u = 20$, the coefficient estimates on the NETPAY_{23} and $P21P3156$ variables are positive and significant. However when $u = 100$, this is no longer the case. Finally, in Table 11(c), we consider whether the marketability variable, $P21P3116$, remains insignificant in logit regressions that use the data from the five sessions of model A where $u = 100$ and the initial distribution of goods over types was made as close as possible to the speculative steady-state equilibrium distribution. Using pooled data from these five sessions (numbers 19–23), we see, once again, that marketability concerns do not seem to be important to subjects even when they start very close to the speculative steady-state distribution of goods over types as evidenced by the insignificant coefficient estimates on the $P21P3116$ variable in all regressions reported in Table 11(c).

It thus appears that player type 1s are not responding to marketability concerns in environments where speculative behavior is their steady-state best response. The only consistently significant determinant of type 1 behavior is the relative past payoff performance of the two strategies, as measured by the reinforcement variable NETPAY_{23}.

B. *Player Types 2 and 3 in Model B*

By contrast with model A, in our parameterization of model B, it is player types 2 and 3 who may rationally choose to speculate in treatments where $u = 500$. Tables 12 and 13 report logit regression results for these two player types using the data from the four sessions involving model B.

Table 12(a) reports regression results using the pooled, type 2 player data from the 2 sessions (numbers 5 and 8) where $u = 20$, and Table 12(b) reports regression results using the pooled type 2 player data from the 2 sessions (numbers 6 and 7) where $u = 500$. When $s^2_{13} = 0$, type 2 players play the fundamental strategy of refusing to offer good 1 for good 3 and similarly, when $s^2_{31} = 0$, type 2 players play the fundamental strategy of offering to trade good 3 for good 1. We see that for both treatments, $u = 20$ and $u = 500$, the coefficient on the reinforcement variable, NETPAY_{13}, is positive and significant in explaining the probability of fundamental play by type 2 players with good 1 who face good 3 in trade, but it is insignificant in explaining the behavior of type 2 players with good 3 who face good 1 in trade. Furthermore, the coefficients on the relevant marketability variables, $P21P3243$ and $P12P32475$ are not significantly different from zero in any of the regressions that include the NETPAY_{13} variable. When the

[18] We note that since very few type 1 players in model A ever trade for good 3, even in the "speculative treatment" where $u = 100$, there are considerably fewer observations for the behavior of type 1 players with good 3 who face good 2 than there are for type 1 players with good 2 who face good 3.

TABLE 12—REGRESSION RESULTS FOR PLAYER TYPE 2S IN MODEL B

Independent variable[a]	(a) $u = 20$				(b) $u = 500$			
	s^2_{13}	s^2_{13}	s^2_{31}	s^2_{31}	s^2_{13}	s^2_{13}	s^2_{31}	s^2_{31}
Constant	1.4226** (0.6708)	2.0169*** (0.6183)	−0.8118 (2.1689)	−1.1702 (2.1302)	0.8342*** (0.3007)	1.8268 (0.2464)	−0.6190 (0.5396)	−0.6613 (0.4959)
NETPAY$_{13}$	0.0888*** (0.0243)	—	0.0377 (0.0571)	—	0.1449*** (0.0325)	—	−0.0245 (0.1242)	—
$P12P3243$	0.6201 (0.6182)	1.2105** (0.5934)	1.7591 (1.7703)	2.2146 (1.6670)				
$P12P32475$					−0.0540 (0.4742)	0.4301 (0.4417)	1.2507 (1.5671)	1.2130 (1.5538)
Pr > χ^2 [b]	0.0001	0.0044	0.4183	0.3072	0.0001	0.5664	0.0166	0.0061
Observations	504	504	45	45	337	337	44	44

Notes: Standard errors are reported in parentheses.
[a] The coefficient estimate on one session dummy has been omitted.
[b] p-values from a Wald test of the hypothesis that the coefficient estimates on the independent variables (other than the intercept) are all zero.
** Significantly different from 0 at the 5-percent level.
*** Significantly different from 0 at the 1-percent level.

TABLE 13—REGRESSION RESULTS FOR PLAYER TYPE 3S IN MODEL B

Independent variable[a]	(a) $u = 20$				(b) $u = 500$			
	s^3_{21}	s^3_{21}	s^3_{12}	s^3_{12}	s^3_{21}	s^3_{21}	s^3_{12}	s^3_{12}
Constant	3.7464** (1.6518)	3.8016** (1.6426)	1.6827 (2.3274)	1.9026 (2.3068)	1.5261*** (0.2318)	0.9213*** (0.1852)	1.8801*** (0.4337)	1.8051*** (0.4086)
NETPAY$_{12}$	0.0335 (0.0348)	—	0.1200** (0.0507)	—	0.1957*** (0.0349)	—	0.0333 (0.0591)	—
$P2312$	−0.4683 (3.3355)	−0.6975 (3.3100)	4.4106 (4.8793)	3.6050 (4.8003)				
$P23150$					−0.2773 (1.6019)	−0.3549 (1.5297)	−3.2437 (2.9779)	−3.1249 (2.9663)
Pr > χ^2 [b]	0.0011	0.0005	0.0003	0.0015	0.0001	0.0004	0.2244	0.1318
Observations	487	487	230	230	361	361	132	132

Notes: Standard errors are reported in parentheses.
[a] The coefficient estimate on one session dummy has been omitted.
[b] p-values from a Wald test of the hypothesis that the coefficient estimates on the independent variables (other than the intercept) are all zero.
** Significantly different from 0 at the 5-percent level.
*** Significantly different from 0 at the 1-percent level.

NETPAY$_{13}$ variable is suppressed, we see that in the logit regression for s^2_{13}, using data from the two sessions where $u = 20$, the coefficient on the marketability variable, $P12P3243$, is positive and significant, indicating that type 2 players are more likely to refuse to trade good 1 for good 3 (to play fundamental) when the marketability variable

$P12P3243$ indicates that fundamental play is their best response.[19] Notice however, that in the sessions where $u = 500$, the coefficient on the marketability variable $P12P32475$ remains insignificant when we eliminate the NETPAY_{13} variable from the regression specification. Thus in environments where player type 2s may rationally choose to play speculative strategies, we see once again that marketability considerations do not appear to be a significant concern of the players.

Table 13(a) reports logit regression results using the pooled, type 3 player data from the 2 sessions (numbers 5 and 8) where $u = 20$, and Table 13(b) reports logit regression results using the pooled type 3 player data from the 2 sessions (numbers 6 and 7) where $u = 500$. Here, $s_{21}^{3} = 0$ when type 3 players play the fundamental strategy of offering good 2 for good 1 and similarly, $s_{12}^{3} = 0$ when type 3 players play the fundamental strategy of refusing to trade good 1 for good 2. We see again that only past relative payoff differences, as measured by the reinforcement variable NETPAY_{12} seem to matter in explaining type 3 behavior across treatments and in different trading situations.

From our results for player type 2s and 3s in model B, we conclude that, like player type 1s in model A, these players are not responding to marketability considerations when choosing trading strategies, and instead appear to be mainly motivated by past relative payoff performance.

[19] In treatments where $u = 20$, the variable $P12P3243$ is, by construction, *always positive;* unlike other marketability variables, it can never be negative. It follows that fundamental play by player type 2s with good 1 facing good 3 is always the unique best response when $u = 20$. Furthermore, the variable NETPAY_{13} is always positive in the two sessions of model B where $u = 20$; this has the effect of reinforcing fundamental play by player type 2s with good 1 who face good 3. Indeed, in the two sessions where $u = 20$, the correlation between $P12P3243$ and NETPAY_{13} is 0.15. In our other treatments, the correlation between the marketability and netpay variables is typically closer to zero. Thus, we usually find that the marketability variable remains insignificant when we purge the netpay variable from our regression specification.

VI. Conclusions

The Kiyotaki-Wright (1989) model environment captures the essential role of media of exchange in overcoming search costs when production is specialized and distribution is decentralized. In this environment, no one can secure their own consumption good unless some individuals are willing to accept in trade a commodity that they do not wish to consume themselves. An agent who accepts in trade a good other than that agent's consumption good must believe that he is in a better position with respect to the future having made the trade than he would have been if the trade were refused. It is these beliefs that give rise to the general acceptability of a commodity as a medium of exchange. Kiyotaki and Wright show that if agents form their beliefs about the expected profitability of holding any given commodity rationally, then an agent must take into account: (i) the per-period cost of holding that commodity, (ii) the utility gained from securing that agent's consumption good, and (iii) the likelihood of meeting someone who is holding that agent's consumption good and is willing to accept the good that agent is currently holding in exchange. The central question addressed in this paper is: do agents who are placed in the Kiyotaki-Wright environment behave as if they were forming their beliefs in this way?

Prior simulations by Marimon et al. (1990) demonstrated that agents who followed an adaptive learning algorithm would endogenously generate media of exchange in the Kiyotaki-Wright environment. However, this population of artificial agents tended to settle into fundamental strategies even when the rational expectations equilibrium required one or more of the types to follow a speculative strategy. That is, agents who are programmed to only respond to past payoffs do not generate the patterns of behavior that would support fully rational beliefs in the Kiyotaki-Wright environment. A similar pattern of behavior was also observed in a behavioral experiment by Brown (1996). However, Brown's experiment does not incorporate all of the features of the Kiyotaki-Wright environment, and does not allow for a comparative static analysis.

Our experimental design was intended to provide as close an approximation to the Kiyotaki-

Wright environment as is possible in a laboratory setting. Several notable features of our design were our consideration of different treatments and model versions, our effort to induce risk neutrality by paying subjects according to a binary lottery, our effort to implement an infinite horizon environment with constant discounting and finally, our effort to implement the common knowledge assumption by informing subjects of the historical average proportions of goods held by each player type in the population. All of these features have been missing from previous analyses. The last two, in particular, would seem to be crucial to agents' ability to achieve coordination on a speculative equilibrium.

We found that the modifications we made to the experimental environment in order to map the theoretical environment as closely as possible had little influence on behavior. Our subjects showed a pronounced tendency to play fundamental strategies regardless of treatment conditions. When subjects did respond to increases in the utility value of consumption by increasing the frequency with which they played speculative strategies, this was often done by agent types who, in theory, ought not to have speculated. The dominance of fundamental strategies was unaffected by our efforts to initialize inventory holdings so that they were close to the speculative equilibrium distribution of goods and was also unaffected by various different information treatments intended to promote speculative play. At the individual level, behavior reflected a response to differences in past payoffs—as assumed in reinforcement learning models, but did not reflect any response to differences in marketability conditions—as required by "fully rational" Bayesian models.

The negative results we find with regard to the equilibrium predictions of the Kiyotaki-Wright (1989) model should be interpreted with care. It may very well be that the way in which this particular model frames the dynamic problem for agents does not make the nature of the trade-off between reducing current storage costs and increasing future expected payoffs particularly transparent. Furthermore, the sharp distinction between the patterns of behavior that are generated by adaptive behavior and those generated by fully rational agents that is observed in this model is not necessarily characteristic of other search-theoretic models of money.[20] Finally, while we have done some robustness checking, we have, as in any experimental study, considered only a few different parameterizations of the Kiyotaki-Wright model and we have not considered all possible information treatments, as time and budget limitations prohibit such an endeavor.

If, on the other hand, our experimental implementation of the Kiyotaki-Wright model incorporates in a sharp and clear manner the essence of the kinds of trade-offs that are inherent in dynamic environments, then our findings have broader implications that extend beyond the search-theoretic framework that we consider. In particular, our results appear to call into question the empirical usefulness of the comparative dynamic implications derived from models where individual decisions are characterized by solutions to dynamic programming problems. Indeed, much remains to be learned about the range of empirically valid propositions that can be gleaned from dynamic optimization models, including search-theoretic models of money.

APPENDIX

This Appendix provides the written instructions that were given to subjects in experimental sessions 2, 3, 9, 11, and 13. The instructions for other sessions are similar. These written instructions were read aloud prior to the beginning of play. In addition, a single practice trading round was played in order to familiarize subjects with the computer interface.

INSTRUCTIONS

General

You are about to participate in an experiment in economic decision-making. Funding for this ex-

[20] See for example, Wright (1995), who presents a model in which the distribution of players over types evolves according to an evolutionary (reinforcement-based) dynamic while at any given point in time exchange behavior is governed by the expected profitability conditions of the model with a fixed distribution of types.

periment has been provided by the National Science Foundation. Please read these instructions carefully. If you have any questions please feel free to ask. We ask that you not talk with one another during the experiment.

This experimental session is divided up into a number of games. Each game consists of a number of rounds.

Participants in this session have been divided up equally into one of three different types: type 1, type 2, or type 3. Your type was chosen randomly at the beginning of the session and will be indicated in the top left corner of your computer screen. Your type will not change for the duration of the session.

In addition to the three different types of players, there are also three different types of goods: good 1, good 2, and good 3.

Each player begins each game with 100 points and one unit of the good that he or she "produces" in storage. Each player type produces a good that is different from his or her type:

Player Type	Produces
1	Good 2
2	Good 3
3	Good 1

Players may continue to hold the good they produce in storage, or they may exchange this good for one unit of the good held in storage by another player. Exchanges are always one for one.

Your Objective

Your objective in every round of every game is to get as many points as possible. You earn a positive number of points only when you obtain the good corresponding to your type. Thus, type 1 players will want to obtain as many units as possible of good 1, type 2 players will want to obtain as many units as possible of good 2, and type 3 players will want to obtain as many units as possible of good 3. Storing goods costs you points. Storing good 1 costs you 1 point per round, storing good 2 costs you 4 points per round, and storing good 3 costs you 9 points per round. The more points you earn the greater is your probability of winning an additional prize of $10 (for a total of $20). Since each player produces a good that is different from his or her type, players must *trade* to get the good that corresponds to their type. Notice that no player of type j produces the good desired by the player type who produces good j. Therefore, to get the good corresponding to your type, you may have to engage in more than one trade.

Each game consists of a sequence of rounds. We will now describe how a round is played.

The Play of A Round

In each round, players of all types are randomly paired with one another. On the first screen that you see, you are told the type of the player with whom you are matched in the current trading round, and the type of good that this other player has in storage. You are also reminded of the good you currently hold in storage. If the other player agrees, you may trade the goods that you both have in storage. You are asked:

Do you want to trade your good for the other player's good? (Y/N):

Type **Y** for Yes, to indicate that you want to trade the good you currently have in storage for the good the other player has in storage or type **N** for No, to indicate that you do not want to make the trade. You will then be asked whether you want to change your decision. Type **Y** only if you want to *change* your decision. Type **N** otherwise.

The Outcome of a Round

Once all players have made their trading decisions, a second screen appears that tells you the outcome of the last round. There are *three* possible outcomes for each trading round:

1. You proposed to trade, your proposal was accepted and you received the good corresponding to your type. In this case, you earn a positive net payoff in points as determined in Table 1 below.

Table 1

Type	Points for Obtaining Good Corresponding to Type		Storage Cost of Good Produced		Net Payoff
1	100	−	4	=	96 points
2	100	−	9	=	91 points
3	100	−	1	=	99 points

You receive 100 points for obtaining the good that corresponds to your type minus the storage cost for storing a unit of the good you produce. Whenever you receive the good corresponding to your type, you immediately produce a new unit of your production good so that at the end of the round, the good you have in storage is your production good.

2. You proposed to trade, your proposal was accepted and you received in trade a good that does not correspond to your type. In this case, you lose a certain number of points corresponding to the storage cost of the good you received from the other player as determined in Table 2 below. The good you now have in storage is the good you received from the other player.

Table 2

Good	Storage Cost Per Round
1	1 point
2	4 points
3	9 points

3. You or the player with whom you were matched chose not to trade. In this case no trade occurs. The good you now have in storage does not change from the previous round. You lose a certain number of points corresponding to the storage cost of the good you hold in storage, as determined by Table 2.

Note that at the end of every round of every game you always have in storage one unit of a good that does not correspond to your type. Therefore, in every round, you always lose a certain number of points corresponding to the storage cost of the good you have in storage as determined by Table 2. You earn a positive net payoff in points, as determined in Table 1, only when you obtain the good that corresponds to your type.

When Does a Game Continue and When Does It End?

After the outcome of the previous round has been revealed to all players, a random number is drawn from the interval 1 to 100. If this random number is less than or equal to 90, the game continues on with another round. If the random number is greater than 90, the just-completed round will be the last round of the game. Thus, after every round there is a one-in-ten chance that the game will not continue into the next round. The bar chart in the middle of the trading screen reflects the probability that the game will end 1–10 rounds from the current round.

When a game ends, you will see a message on your screen. You will be told your point total for that game. Then, depending on the time available, a new game will begin. You will start the new game with 100 points and the good you produce in storage.

Strategic Considerations

Before making your trading decisions, you may want to take account of some of the information that is available to you on the first screen that you see.

At the top of this screen, you are reminded of your type, the round number, and your new point total for the current game, as of the last trading round. You are also reminded that you receive 100 points for the good corresponding to your type, and the costs in points for storing a unit of each of the three goods in every round. You may want to take these storage costs into account when deciding whether or not to trade the good you currently have in storage.

In the middle of your screen is a bar chart indicating the cumulative probability that the game will end 1 to 10 rounds from the current round. This chart reflects the 1 in 10 chance that the game will end from one round to the next. Observe that this probability is increasing, indicating that it is increasingly likely the game will end 1 to 10 rounds from the current round. Since the game may end soon, you may not be able to meet with a player who is willing to trade for the good you currently have in storage. Therefore, in addition to considering storage costs, you may also want to consider the *time* it will take you to trade for the good that corresponds to your type, when choosing whether or not to trade the good you currently have in storage.

At the bottom of your screen you will see a table listing the percent of each type of player that is storing each type of good. These percentages are average percentages over all rounds that have been played in the current game, and are updated at the end of every trading round. You may want to use this information in making an estimate of how long it will take you to meet a player who both has the good corresponding to your type, and who will want to trade it for the good you currently have in storage, or for some other good you could get in trade.

Earnings

All subjects who complete this 2-hour experimental session are guaranteed to receive a $10 payment. Depending on how many points you earn, it is possible for you to earn an additional $10 prize for a total of $20.

Following the last game of the session, your point total from *one* game, chosen at random from all of the games that you played, will be converted into a probability of winning the additional $10 prize. You will have a positive probability of winning the $10 prize if your point total from the game chosen exceeds the initial total of 100 points that you are given at the beginning of every game. If your point total in a game falls below the initial 100 point total, your probability of winning the $10 prize is 0 if that game is the one chosen at random.

Your probability of winning the $10 prize depends only on how many additional points you were able to obtain in the game chosen relative to the maximum number of points that were obtainable in this same game *for a player of your type*. You are not competing with other players for the $10 prize. Each of you could win the $10 prize if you earned enough additional points in the game chosen. Note that each additional point you earn over the 100 initial points you are given increases your probability of winning the $10 prize by same amount. Thus, the more points you earn in a game over the initial point total of 100, the greater is your probability of winning the $10 prize, if that game is chosen at random. Since you do not know at the outset of play what game will be chosen at random, your objective in every game should be to obtain as many points as is possible.

ARE THERE ANY QUESTIONS BEFORE WE BEGIN?

REFERENCES

Aiyagari, S. Rao and Wallace, Neil. "Existence of Steady States with Positive Consumption in the Kiyotaki-Wright Model." *Review of Economic Studies*, October 1991, *58*(5), pp. 901–16.

———. "Fiat Money in the Kiyotaki-Wright Model." *Economic Theory*, October 1992, *2*(4), pp. 447–64.

Brown, Paul M. "Experimental Evidence on Money as a Medium of Exchange." *Journal of Economic Dynamics and Control*, April 1996, *20*(4), pp. 583–600.

Duffy, John and Ochs, Jack. "Fiat Money as a Medium of Exchange: Experimental Evidence." Working paper, University of Pittsburgh, 1999.

Holland, John H. "Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems," in R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, eds., *Machine learning: An artificial intelligence approach*, Vol. 2. Los Altos, CA: Morgan Kaufmann, 1986, pp. 593–623.

Kehoe, Timothy J.; Kiyotaki, Nobuhiro and Wright, Randall. "More on Money as a Medium of Exchange." *Economic Theory*, April 1993, *3*(2), pp. 297–314.

Kiyotaki, Nobuhiro and Wright, Randall. "On Money as a Medium of Exchange." *Journal of Political Economy*, August 1989, *97*(4), pp. 927–54.

———. "A Contribution to the Pure Theory of Money." *Journal of Economic Theory*, April 1991, *53*(2), pp. 215–35.

———. "A Search-Theoretic Approach to

Monetary Economics." *American Economic Review*, March 1993, *83*(1), pp. 63–77.

Marimon, Ramon; McGrattan, Ellen and Sargent, Thomas J. "Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents." *Journal of Economic Dynamics and Control*, May 1990, *14*(2), pp. 329–73.

Roth, Alvin E. and Malouf, Michael W. K. "Game-Theoretic Models and the Role of Information in Bargaining." *Psychological Review*, November 1979, *86*(6), pp. 574–94.

Sethi, Rajiv. "Evolutionary Stability and Media of Exchange." Working paper, Barnard College, Columbia University, 1996.

Siegel, Sidney and Castellan, N. John, Jr. *Nonparametric statistics for the behavioral sciences*, 2nd Ed. New York: McGraw-Hill, 1988.

Wright, Randall. "Search, Evolution and Money." *Journal of Economic Dynamics and Control*, January/February 1995, *19*(1–2), pp. 181–206.