Intrinsically Worthless Objects as Media of Exchange: Experimental Evidence*

John Duffy
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260 U.S.A.
Tel: (412) 648–1733
FAX: (412) 648–1793
E–mail: jduffy@pitt.edu

Jack Ochs
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260 U.S.A.
Tel: (412) 648–7061
FAX: (412) 648–1793
E–mail: jochs@pitt.edu

This Draft: December 2000

Abstract

This paper reports results from an experimental implementation of Kiyotaki and Wright's (1989) search–theoretic model of money as a medium of exchange. The focus of the experiment is on the role played by 'token' objects in this environment. A token is an object that serves as a store of value, but which has no intrinsic value to any individual. The question we address is whether such token objects may also serve as *media of exchange*. According to the theory, there are parameterizations of the Kiyotaki–Wright (1989) model under which a single token object serves as a medium of exchange in a pure strategy Nash equilibrium. Moreover, as Aiyagari and Wallace (1992) have shown, this token object need not be the least costly–to–store good. Our experiment implements the Kiyotaki–Wright (1989) environment with a single token object and tests whether these equilibrium predictions are robust to the dynamics created by out–of–equilibrium behavior. Our findings suggest that there is some support for the theoretical predictions; in particular, we find that subjects nearly always offer to trade for the token object when such a trade lowers their storage costs. However, contrary to the theory, we also find that subjects frequently refuse to offer to trade the token object for more costly–to–store goods when the theory predicts they should make such trades.

Journal of Economic Literature Classification Nos. C92, D83, E40.

*Financial support for this project was provided by the National Science Foundation under grant # SBR95–07968. We thank the two anonymous referees and seminar participants for their insightful comments and constructive suggestions. We are also grateful to Sridhar Komandur and Robert Conticello for their assistance in developing the program used to carry out the experiments. The raw data from all experiments reported in this paper are available at: `www.pitt.edu/~jduffy/fiat/`

1 Introduction

By definition, money is an object that serves as a medium of exchange. As Ostroy and Starr (1990) have emphasized, a medium of exchange must also be a store of value, but not all stores of value are media of exchange. The defining characteristic of an object that is both a store of value and serves as a medium of exchange is its *general acceptability* to individuals when offered. Consequently, a storable object that becomes a medium of exchange need not have value to anyone as either a consumption good or as an input into a physical production process. Of course, an individual who accepts an object that has no value to that individual as either a consumption good or as an input to a production process does so in the belief that the object will be acceptable in exchange for something of direct utility to that individual.

A storable object with no intrinsic value in consumption or production is a "token". Governments may make tokens media of exchange by fiat, requiring that these objects be accepted for the settlement of all debts public and private. However, when a government destroys the storage value of a fiat object by excessive issue, that object will cease to serve as a medium of exchange, regardless of the government's declaration.

According to the theory propounded by Kiyotaki and Wright (1989), it is not important that an object have either intrinsic value in use, or value created by a government's fiat declaration, in order that the object serve as a medium of exchange. It is important that the object be durable and not too costly to store. But in their theory, an object with no other intrinsically valuable characteristics can acquire and maintain the status of a medium of exchange simply by virtue of the fact that people come to expect it to be generally accepted in exchange for objects that have intrinsic value. That is, an object becomes a medium of exchange because people have formed a (rational) expectation that it will continue to serve as a medium of exchange and this expectation need not be supported by any property of the object other than the social convention that has emerged from its use in the past.

In prior work (Duffy and Ochs (1999)) we reported the results of an experiment designed to test the implications of one version of Kiyotaki and Wright's (1989) model, where every object had (intrinsic) consumption value to one player type, but no object had consumption value to all player types. In that version of the Kiyotaki–Wright model (as well as in the version studied in this paper), each player could store only one unit of any object in every period. Players were randomly matched pairwise, and could engage in one for one trades of the goods they held in storage. If a player successfully traded for his "consumption good," he immediately produced a unit of his "production good", so that no player was ever storing the object that yielded him consumption value. The storage costs varied across the objects. In the environment studied by Duffy and Ochs (1999), each individual was frequently offered a good that had no intrinsic value as a consumption good to him/herself, though it did have consumption value to a certain fraction of the population. When faced with the opportunity to trade for such a "non-consumption" good, an individual could exhibit either of two behaviors: fundamental, or speculative. Fundamental behavior is exhibited when an individual accepts the offer if and only if the offered good has a lower storage cost than the good the individual is currently storing. Speculative behavior is exhibited if the individual accepts such a good even if it has a higher storage cost than the good that is currently held. Speculation involves accepting a good in anticipation that the higher storage cost good will allow the holder to execute a successful future trade for a good of intrinsic value to the holder sooner than the good that the individual currently gives in exchange. The theory predicts that under some parameter values only fundamental behavior will be observed in equilibrium and, as a consequence, the good with the lowest storage cost will emerge as a generally accepted medium of exchange. Of particular interest, the theory also predicts that under some parameter values speculative behavior will be observed and that this behavior will support an equilibrium in which a good that does not have the minimum storage cost will emerge as a medium of exchange. Our prior experiment concluded that the theory was partially successful in accounting for the emergence of specific commodities as media of exchange. When the environment had parameter values for which the theory predicted that the object with the lowest storage costs would emerge as the unique medium of exchange, behavior tended to conform to this prediction. However, when the environment had parameter values for which the theory predicted that an object with more than the minimum storage costs would emerge as a medium of exchange, the observed behavior indicated that speculative behavior was observed much less frequently than predicted and no object which had less than the minimum storage cost emerged as a generally accepted medium of exchange. This finding was consistent with the report of an experiment conducted by Brown (1996), who used a somewhat different experimental design, and with a simulation study of artificially intelligent agents who followed an adaptive learning process that was conducted by Marimon et al. (1989).

Can a good that has *no* value as a consumption good to anyone, but which is less costly to store than any other durable good emerge as a medium of exchange, as the theory of Kiyotaki and Wright implies? The fact that the lowest cost good emerged as a medium of exchange in our prior experiment suggests that a good with no consumption value, but also with the lowest cost of storage, may, in fact, emerge as a medium of exchange. On the other hand, accepting a token object with *no value* as a consumption good to any agent *is inherently an act of speculation*. The observed reluctance of many people in the prior experiments of Duffy and Ochs (1999) and of Brown (1996) to engage in speculative behavior, suggests that a good with no intrinsic value may not emerge as a medium of exchange, even if it has the lowest storage costs. The experiment reported below was designed to answer this question and, in doing so, to provide a cleaner test of the Kiyotaki-Wright theory.

Specifically, in the experiment reported below, a fourth object is added to the environment of the previous experiment. This object, good 0, occupies storage space, is in fixed supply, and cannot be destroyed. It has no consumption value to anyone. Some individuals are always holding good 0. Since good 0 is not a consumption good for anyone, if it is to be accepted in exchange, the person who accepts it must be speculating that someone else will value good 0 for its use in further exchanges, and that belief implies the further belief that yet other agents will wish to accept the object to facilitate further exchanges, ad infinitum. This belief in the general acceptability of an object is the true essence of a medium of exchange and, according to the Kiyotaki-Wright theory, can become attached to any storable object, provided that its cost of storage is sufficiently low relative to other objects that might serve as media of exchange. Indeed, an extension of the Kiyotaki-Wright model provided by Aiyagari and Wallace (1992), a version of which is also examined in this paper, predicts that under certain parameterizations, an intrinsically worthless object will have some market acceptability even when it is not the least costly-to-store object. To foreshadow the results of our experiment, we find that when good 0 is the good with the lowest storage cost, and the parameters are set so as to predict that the lowest storage cost good will emerge as a medium of exchange, good 0 is, in fact, as widely used as a medium of exchange as was the intrinsically valued good with lowest storage in our previous experiment. In addition, even when good 0 is not the least costly–to–store good, it will circulate. However, in this case, the degree of acceptability of good 0 and its consequent rate of circulation is significantly less than the theory predicts.[1]

[1]There have been several other experimental studies involving token objects as a *store of value*, e.g. Bernasconi and Kirchkamp (2000), Lian and Plott (1998), Lim, Prescott and Sunder (1994), Marimon and Sunder (1993, 1994) and McCabe (1989); Duffy (1998) surveys this literature. Unlike these studies, the Kiyotaki and Wright (1989)

2 The Environment

We study versions of Kiyotaki and Wright's (1989) search–theoretic model with a single token object. A finite population of N agents is divided up equally into one of three types, referred to as type $\tau = 1, 2, 3$.[2] There are four different indivisible goods, indexed by $j = 0, 1, 2, 3$. Good 0 is the token object; no agent desires to consume this good, and for this reason, it is intrinsically worthless. The supply of good 0 is exogenous. Type τ desires to consume good $j = \tau$ but produces good $\tau + 1$ modulo 3.[3] Thus, type 1 desires good 1 but produces good 2, type 2 desires good 2 but produces good 3, and type 3 desires good 3 but produces good 1. Notice that in this environment there is an absence of what Jevons termed a "double–coincidence–of–wants" that can only be overcome if some agents are willing to trade for goods they do not value, but which they expect they will be able to later trade for goods they do value. Such goods are regarded as media of exchange.

Agents have access to a storage technology that enables them to store one unit of any good in every period. The fraction of the population storing good 0 is an exogenously given constant, $m \in (0, 1)$. The remaining fraction of the population, $1 - m$, have intrinsically valued goods in storage. Storage, however, may be costly. The per period cost (in terms of utility) of storing a unit of good $j = 0, 1, 2, 3$ is denoted by $c_j \geq 0$. We assume that these four storage costs are common across player types. Similarly, for simplicity, we assume that all 3 player types receive the same utility from consumption $u > c_3$ and have the same discount factor $\beta \in [0, 1]$.

At the beginning of a trading period all N agents are randomly paired with all possible pairings equally likely. Each agent must then decide whether to trade the good that he currently holds in storage for the good that the player with whom he is matched currently holds in storage. A trade occurs only if both parties agree to trade. Because goods are indivisible and agents can only store a single unit of any good, exchanges always involve one–for–one swaps of goods in storage.[4]

environment we examine is one in which multiple, durable and intrinsically valued goods compete with the unique token object as stores of value, thus providing a more rigorous test of whether a token object may serve as a medium of exchange.

[2]Kiyotaki and Wright (1989) assume there is a continuum of agents, however in our experimental design, we are constrained to considering a finite population of size N. We have made an effort to consider population sizes of $N = 18$ or $N = 24$, that are large by comparison with other experimental studies, in an effort to obtain a good approximation to the continuum of agents that the theory presumes. Appendix 1 discusses the consequences of our finite population approximation.

[3]Kiyotaki and Wright refer to this version of their model as "Model A". In another version, "Model B", which is not isomorphic, type τ produces good $\tau - 1$ modulo 3. See also Aiyagari and Wallace (1991, 1992) who generalize the model to allow for n types and n goods plus a token object.

[4]See however, Shi (1995) and Trejos and Wright (1995) who relax the indivisibility assumption thereby permitting prices to be endogenously determined. Such a design lies beyond the scope of this paper.

If a player of type τ successfully trades for good τ, he consumes that good, receiving utility u, and he immediately produces a new unit of his production good, $\tau + 1$ modulo 3, which becomes the good he holds in storage. Furthermore the player is immediately charged the storage cost of his production good so that his net payoff from successfully trading for his consumption good τ is $u - c_j > 0$, where $j = \tau + 1$ modulo 3. If a player of type τ successfully trades for some good $j \neq \tau$, the good he has in storage is changed to reflect this trading decision. If a trade is not mutually agreed upon, both players continue to hold the same goods in storage that they held prior to being paired. In either of these last two cases, each player is charged the storage cost of the good he held in storage following the completion of trade, so that his net payoff for the period is $-c_j$, where j is the good held in storage at the end of the round. Since no type produces or consumes good 0, and good 0 cannot be discarded, the total stock of good 0 in the economy, Nm, does not change from period to period, while the total stock of the other goods may vary from period to period, depending on the pairing of agents and their trading decisions.

2.1 Model Parameterization

The storage costs for the three intrinsically valued goods always satisfy:

$$0 \leq c_1 < c_2 < c_3 < u.$$

As our focus is on the use of the intrinsically worthless object as a medium of exchange, we will consider two different cases for the storage cost of this object. The first case is the one studied by Kiyotaki and Wright (1989).

$$\text{Case 1:} \quad 0 = c_0 < c_1$$

In this case, the token object is both costless to store and the least costly–to–store object. The second case proposed by Aiyagari and Wallace (1992).

$$\text{Case 2:} \quad 0 = c_1 < c_0 < c_2.$$

In this case, good 0 is *neither* costless nor the least costly–to–store object. The particular versions of the Kiyotaki–Wright (1989) model we study comprise two main parameterizations, parameter sets 1 and 2, that are consistent with cases 1 and 2 respectfully. These two parameter sets are provided in Table 1. Notice that the value of m varies only *within* parameter set 1. and that storage costs, c_j, the utility value of consumption u, and the value of m all change as we move

Table 1: Parameter Sets Used in the Experiments

Parameter Set	c_0	c_1	c_2	c_3	u	β	m
1	0	1	4	9	20	0.90	0.250, 0.333, or 0.500
2	1	0	23	24	100	0.90	0.167

from parameter set 1 to parameter set 2. These parameter variations were necessary to ensure the existence of the two different stationary Nash equilibria we examine in this paper.

2.2 Steady State Behavior

Our aim is to determine whether subjects adopt trading strategies consistent with the model's steady state predictions. We now discuss these steady state predictions. A more complete analysis is provided in Appendix 1. As our focus is on steady state behavior, we drop all references to time.

Let $v^j_{\tau,i}$ denote the end–of–period, expected discounted utility to player i of type τ from storing good j. This value is explicitly defined in Appendix 1. Let $s^{jk}_{\tau,i} \in [0,1]$ denote the strategy played by type τ player i when storing good j and randomly paired with a player storing good $k \neq j$. If $s^{jk}_{\tau,i} = 1$, the type τ player offers to trade good j for good k; if $s^{jk}_{\tau,i} = 0$ the type τ player refuses to trade.[5] In the steady state, all players of type τ, $i = 1, 2, ..., N/3$ play according to the same, time–invariant strategy; we therefore drop the i subscript. This strategy is optimal in the sense that:

$$s^{jk}_{\tau} \in \begin{cases} \{0\} & \text{if } v^j_\tau > v^k_\tau, \\ \{1\} & \text{if } v^j_\tau < v^k_\tau, \\ [0,1] & \text{otherwise.} \end{cases}$$

As in Kiyotaki and Wright (1989), our focus is on steady state, *pure strategy* equilibria, so $s^{jk}_\tau \in \{0,1\}$, $\forall \tau$, j, k.

As discussed in Appendix 1, under both parameter sets 1 and 2, it is always optimal for player type τ to trade for good τ whenever he meets this good in trade as $v^\tau_\tau = \max_j v^j_\tau$ for all τ. The ranking of the other values depend on the particular parameter set used. Since players always trade for their consumption good and strategies are assumed to be symmetric we can write the stationary pure strategy vector of a type τ player compactly as

$$s_\tau = \left(s^{\tau+1,0}_\tau, s^{k,0}_\tau, s^{\tau+1,k}_\tau \right),$$

[5]If type τ player i with good j meets another player storing good j, we assume that no trade occurs.

where $\tau + 1$ is type τ's production good and $k \neq 0$ is the other intrinsically valued good that type τ neither produces nor consumes.[6]

Under parameter set 1, where good 0 is the least costly–to–store good, Kiyotaki and Wright (1989) show that there exist two stationary pure strategy equilibria. In Appendix 1 we verify the existence of these equilibria for our laboratory implementation of the model. In the first of these, good 0 has positive value, and the strategy profiles of the three player types are:

$$s_1 = (1, 1, 0), \quad s_2 = (1, 1, 1) \quad s_3 = (1, 1, 0)$$

In this steady state, all three types find it optimal to trade for good 0 whenever they encounter it in trade, making good 0 a generally accepted medium of exchange. Furthermore, in this steady state, all players adhere to "fundamental," storage cost–reducing strategies. In particular, type 1 and type 3 players refuse to trade their production goods 2 and 1 for the more costly–to–store goods 3 and 2, respectively, and type 2 players always offer to trade their production good 3 for the less costly–to–store good 1. Hence, good 1 also serves as a (limited) medium of exchange in this steady state equilibrium.

In the other steady state equilibrium, good 0 fails to circulate as a medium of exchange simply because *no one believes it has value in exchange*, i.e. $v_\tau^0 = 0$ for all τ. In this equilibrium, all players refuse to accept good 0 in trade but always offer to accept any other good that has a lower storage cost than the good they are currently holding. In particular, the strategy profiles of the three player types in this steady state are:

$$s_1 = (0, 0, 0), \quad s_2 = (0, 0, 1) \quad s_3 = (0, 0, 0)$$

In this steady state, good 1 is the only medium of exchange.

The frequencies with which players *offer* to trade for certain goods are not the only theoretical predictions that can be tested. It is also possible to test the model's predictions concerning the steady state *probability of acceptance* of goods offered in trade as a function of the proportion of the population that is holding the token object, m. As Kiyotaki and Wright (1989) observe, these acceptance probabilities are a better measure of "moneyness" than more traditional measures such as the velocity of circulation of goods since in this model, goods that are not used as media of exchange (e.g. high storage cost goods) can nevertheless have a very high velocity of circulation in equilibrium, especially as the proportion of the population holding the token object, m, increases.

[6] By "symmetric", we mean that $s_\tau^{kj} = 1 - s_\tau^{jk}$.

The acceptance probability of each good is a function of m because the calculation of this probability uses the steady state proportions of each player type storing each type of good, and these proportions depend on the value of m. The steady state proportions for the three different values that we use for m in parameter set 1 are given in Appendix 1. Using the proportions for the steady state in which good 0 serves as a medium of exchange, one can calculate the number of times each good gets offered in trade per period and the number of times each good is actually traded per period.[7] The steady state probability that a good is accepted when offered is the ratio of the number of times it is traded to the number of times it is offered in each period. These steady state acceptance probabilities are given in Table 2 below.

Table 2: Steady State Probability that a Good Offered in Trade is Accepted.
Case 1 (Parameter Set 1) where Good 0 is Used as a Medium of Exchange

	Acceptance Probability of			
$m =$	Good 0	Good 1	Good 2	Good 3
0.25	1.00	0.576	0.317	0.391
0.33	1.00	0.521	0.313	0.388
0.50	1.00	0.450	0.308	0.378

Since our focus here is on the case where good 0 serves as a medium of exchange, we see that under parameter set 1, the probability that good 0 is accepted when offered is always 100% regardless of the value of m. More importantly, notice that the steady state acceptance probability for good 1, the only other medium of exchange, is also high, but that this acceptance probability decreases as m increases. We note that the acceptance probabilities for goods 2 and 3 also decrease slightly as m increases. However, neither of these goods serve as media of exchange in a stationary equilibrium. We will therefore focus our attention on the hypotheses that good 0's acceptability remains unchanged as m increases, and that good 1's acceptability decreases as m increases.

We next consider the case where good 0 is not the least costly–to–store good, as studied by Aiyagari and Wallace (1992). Under our parameter set 2, we explain in Appendix 1 that there again exist two stationary pure strategy equilibria. In the first, good 0 serves as a limited medium of exchange, even though it is no longer costless to store, and it is more costly to store than good

[7]This can also be done for the other equilibrium, where good 0 is not accepted in trade.

1. The strategy profiles of the three player types in this steady state are:

$$s_1 = (1, 0, 1), \quad s_2 = (1, 0, 1) \quad s_3 = (0, 1, 0)$$

Notice that while type 1 players always offer their produced good 2 for the less costly–to–store good 0, in this steady state, they also find it optimal to speculate in good 3, refusing to trade good 3 for the less costly good 0. Symmetrically, they always offer to trade good 0 for good 3. In fact, a type 1 player will be unable to trade good 0 directly for the desired consumption good 1 without first trading that good 0 for good 3.[8] Type 1 players also play the speculative strategy of trading their produced good 2 for the more costly–to–store good 3.[9]. The steady state strategy profile for type 1 players is the most interesting prediction of the Case 2 environment, and we focus much attention on the behavior of type 1 players in this environment later in section 4.2.

Type 2 and 3 players adhere to fundamental trading strategies as in Case 1; However, the fact that good 1 is now less costly to store than good 0, alters their trading behavior a little. In particular, type 2 players will continue to offer to trade their produced good 3 for good 0, but will refuse to offer to trade good 1 for good 0, as good 1 is now costless to store. For the same reason, type 3 players now refuse to offer to trade their produced good 1 for good 0. Type 2 players continue to offer to trade their produced good 3 for good 1.

Notice that in this steady state there are *three* media of exchange. Good 1, the least costly–to–store good, is a generally accepted medium of exchange. Furthermore, goods 0 and 3 also serve as limited media of exchange; both types 1 and 2 will accept good 0 from one another; type 1 players will accept good 0 in exchange for good 2, and type 2 players will accept good 0 in exchange for good 3. However, type 2 players will refuse to accept good 0 in exchange for good 1. Type 3 players will never accept anything other than good 3 in trade, so good 0 is no longer a generally accepted medium of exchange. Similarly, good 3 serves as a limited medium of exchange in this equilibrium, as type 1 players will accept good 3 from type 2 players in exchange for good 0 or 2, but type 2 players with good 0 or 1 in storage will not agree to trade for good 3.

The other pure strategy steady state in Case 2 is one where good 0 has no value to any type and is therefore refused in all trades. The strategy profiles of the three player types in this steady

[8]They will be unsuccessful trading good 0 to type 2 or type 3 players storing good 1, as both of these player types will always refuse such trades, as discussed below.

[9]As discussed in Appendix 1, in this steady state, the values to type 1 of the four goods satisfy $v_1^2 < v_1^0 < v_1^3 < v_1^1$

state are:

$$s_1 = (0,0,1), \quad s_2 = (0,0,1) \quad s_3 = (0,0,0)$$

This strategy profile is the same as the "speculative" steady state equilibrium described in Kiyotaki and Wright (1989) for the case of three intrinsically valued (commodity) objects and no token object.

The numerical example presented in Aiyagari and Wallace (1992) is based on a parameterization of the model in which m, the fraction of the total population holding good 0, is .0258. This fraction is too small for our laboratory implementation with a finite number of agents. We have therefore constructed another numerical example of a pure strategy equilibrium along the lines of Aiyagari and Wallace's example, where $m = 1/6$.[10] The other parameter choices for this second case are given in parameter set number 2. Following the same procedures outlined in Aiyagari and Wallace (1992) we verified that our parameter choices in parameter set 2 give rise to the two pure strategy, steady state equilibria described above, even in the finite population case. A more detailed description of these calculations is provided in Appendix 1.

3 Experimental Design

The set of experiments we report in this paper can be grouped into two main treatments that make use of either parameter set 1 or 2. A total of 14 experimental sessions were conducted. In the first 9 sessions, labeled A1–A3, B1–B3, and C1–C3, we used parameter set 1 and varied m, the fraction of the population storing good 0. In sessions A1–A3, we set $m = 1/4$, in sessions B1–B3, we set $m = 1/3$ and in sessions C1–C3 we set $m = 1/2$. In the remaining 5 sessions, labeled D1–D5, we used parameter set 2 and set $m = 1/6$. Further characteristics of these sessions are provided in Table 3 and are discussed later in the text.

The 14 sessions reported in this paper were conducted using the University of Pittsburgh's experimental economics computer laboratory. For each session, a group of 18 or 24 subjects (as indicated in Table 3) was recruited. In choosing population sizes we were constrained by the number of computer workstations available in the laboratory and by the requirement that population sizes be multiples of 6 as we wanted to have equal numbers of each of the three types and we wanted all players to be paired with one another in each round of the game. Subjects were primarily University

[10]Other choices for m are possible. The choice of m must be rather small for there to be a pure strategy Nash equilibrium where the token object is used as a medium of exchange and is the second least costly–to–store object. The restrictions on the value for m depend on the other parameters that are chosen for the model.

Table 3: Characteristics of Experimental Sessions

Session Number	Utility Value of Consumption	Equilibrium Type[a]	Number of Subjects	Percentage with Token Object (m)	Number of Games	Number of Rounds[b]
A1	20	F	24	.250	11	102
A2	20	F	24	.250	6	108
A3	20	F	24	.250	6	120
B1	20	F	18	.333	11	110
B2	20	F	18	.333	11	102
B3	20	F	18	.333	11	104
C1	20	F	24	.500	10	107
C2	20	F	24	.500	13	92
C3	20	F	18	.500	12	98
D1	100	S	18	.167	7	67
D2	100	S	18	.167	17	107
D3	100	S	18	.167	8	86
D4	100	S	18	.167	8	99
D5[c]	100	S	18	.167	7	99

[a]F=Fundamental, S=Speculative
[b]Total number of trading rounds from all games played.
[c]Subjects in this session were provided with additional information concerning the population–wide distribution of strategies played.

of Pittsburgh undergraduates. Each subject was recruited for a single 90 minute session. No subject participated in more than one session and no subject had any prior experience with the game.

At the beginning of each session, subjects were randomly assigned to play the role of one of the three player types. They were instructed that their type would not change for the duration of the session. Subjects received written instructions concerning the game they were about to play. A copy of the instructions used in the "D" sessions is provided in Appendix 2 (other instructions are similar). The instructions were the same for all players and this fact was reinforced by reading the instructions aloud. Thus, players were informed of the objectives of each of the three player types. In particular, they learned that each type desired the good corresponding to their own type and that no type desired to consume good 0.

After the written instructions were read aloud and questions were answered a single practice game was played to clarify the instructions and to allow subjects to become familiar with the computer interface. Following the practice round, subjects began playing for money, providing responses on their computer terminals 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. It also collects, updates and reveals to subjects information on the historical average proportions of each type storing each of the four different goods. This information serves as a sufficient statistic for the historical distribution of strategies that have been played.

The 14 experimental sessions involved an average of 10 *games* each, and each game consisted of an average of 10 *trading rounds*. Thus, each session involved an average of 100 trading rounds. The total number of games and trading rounds for each session is reported in Table 3. Each trading round can be described as follows.

At the start of the trading round the computer program randomly pairs all subjects in the session. Once paired, the first computer screen on each player's monitor reveals information about the player with whom they are matched – i.e. the type of this other player and the good this other player has in storage – see Figure 1. Players are also reminded of their *own* type, the good they currently have in storage, the storage costs in points of the four goods, 0, 1, 2 and 3, and the number of points they earn when they obtain the good corresponding to their type. Over the course of a game players accumulate points rather than cash balances. Each player's point total for the current game, as of the last trading round is indicated on this first trading screen. The manner in which points are converted into cash payments is discussed below.

Players are asked whether they want to trade the good they have in storage for the good of the player with whom they are matched. They respond to this question by typing either Y for yes or N for no and they have an opportunity to change their decision before it is finalized. Trade occurs if and only if both parties agree to a trade; all trades consist of one–for–one exchanges of the two goods the paired players have in storage.

[Figure 1 here.]

Before making their decisions, players could consider some additional information that was provided on the first trading round screen – refer again to 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 cumulative probabilities in this bar chart reflect the constant, one–in–ten chance (.10 probability) that the game will end from any one round to the next. We chose to have a random stopping rule for each game for two reasons. First, the Kiyotaki–Wright model envisions that agents are infinitely–lived. Second, the model also specifies that agents discount the future using a constant discount factor $\beta \in (0,1)$ that is common across types. The random stopping rule provides us with a mechanism for implementing both of these features of the Kiyotaki–Wright environment. As indicated in Table 1, β is always set to .90 in both parameter sets 1 and 2. To implement discounting of future payoffs by 90%, we chose to end each game with probability $1-\beta$, i.e. with probability 10% at the end of any round that has been reached. Following the completion of a trading round, the computer program draws a random number from a uniform distribution over $[0,1]$. If this random draw is less than or equal to .90, the game continues on with another round. Otherwise, the just completed round is declared to be the last trading round of the game. Subjects were told that there was a constant, one–in–ten chance that each game would end from one round to the next. Thus, the expected length of each game was $\frac{1}{1-\beta} = 10$ rounds. To avoid overly long games, we chose to end any game that exceeded 40 trading rounds. Subjects were not informed that this upper bound of 40 rounds was in effect; they were only told of the .10 probability that the game would end from one round to the next. We note that the upper bound of 40 rounds was reached in just 2 of the 138 games played in all 14 experimental sessions.

At the bottom of the first trading round screen subjects were presented with a table indicating the percentage of each type of agent storing each type of good as of the end of the last trading round. These percentages are *average* percentages based on the entire history of the current game up through the last trading round and are updated following the completion of every trading round. Subjects were informed of this information and how it is calculated in the instructions. They were further instructed that they might want to use this information in forming an estimate of how long it might take them to meet a player who both had the good corresponding to their type and who would be willing to offer this good in trade (see the instructions in Appendix 2). The information on the proportion of each type storing each good is assumed to be common knowledge in the theoretical Kiyotaki–Wright environment. Revealing this information to subjects was our way of implementing this common knowledge assumption.

Our implementation of discounting and the information we provide players about the inventory distributions of goods by type are two features which differentiate this study and our earlier study, Duffy and Ochs (1999), from previous studies of adaptive behavior in the Kiyotaki–Wright (1989) model. (See, e.g. Marimon et al. (1989) and Brown (1996)). We believe that these two features of our design should facilitate the play of speculative strategies (in environments where speculation makes sense) in that they encourage subjects to consider not only the storage cost of the four goods, but also the relative likelihood that they will be able to meet a player who both has the good they desire and who is willing to engage in trade. Furthermore, information concerning the proportion of individuals storing good 0 may facilitate the adoption of this token object as a medium of exchange.

After all players' trading decisions have been made, any mutually agreed upon trades are implemented and the results of each pair's trading decisions are revealed to subjects on a second trading round screen which is also illustrated in Figure 1. At the bottom of this second screen, the information on the economy–wide historical average percentage of each type of agent storing each type of good has been updated to take account of what has occurred in the just completed round of play.

Once the results of a trading round have been made known to all players, the computer draws a random number to determine whether the game will continue on with another round. If the random draw is such that the game continues (which occurs with probability .90), players see a message on their second trading round screen indicating that the game will continue. In this case, they carry the goods they have in storage over to the new trading round. Otherwise, (with probability .10) the just completed round is declared to be the last round of the game, and players see a message at the bottom of the second trading round screen indicating that the game has just ended.

If the game continues, players are randomly paired once again and are asked to make another trading decision. If the game has ended, the players' end–of–game point total is stored by the computer and, depending on the time available, a new game may begin.

Players begin every new game (including the first) with a total of 100 points. Points are then added or subtracted from each player's point totals in accordance with their trading decisions. The good that each player has in storage at the start of every new game depends on their type. A certain fraction, $1 - m$, of each type begins a new game with the good their type *produces* in storage. The remaining fraction, m, of each type begins a new game with good 0 in storage.[11] The decision as to whether a player begins a new game with their production good or good zero in storage is random.

[11]In Duffy and Ochs (1999) we considered two different schemes for the initial distribution of goods over types. The first was similar to the one used here, in that each type began each game with their production good in storage (there were no intrinsically worthless token objects). In the second scheme, the initial distribution of goods over types was made as close as possible to the steady state Nash equilibrium distribution. We found that there was no significant difference in subject behavior from using either one of these two initialization schemes. Therefore, we chose to consider only one initialization scheme in the experiment reported here–the one in which each type begins each game with either their production good or good 0 in storage.

Since there are equal numbers of each player type, and since good 0 is not produced, consumed or discarded, the fraction m also represents the fraction of the entire subject pool that is storing good 0 at any moment in time. Our parameter choices for m are given in Table 1.

Note that with this initialization scheme, players may start each new game with 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 percentage of each type of player storing each type of goods is reinitialized to reflect the initial distribution of goods by type in the new game. Thus, with each new game, some of the players' available information changes. However, since a player's type remains unchanged over an entire session as does the parameterization of the game itself and the pool of subject participants, each player can draw upon his or her own past experience when making trading decisions in the new game.

The total number of games that were played was determined by us according to the time available. Subjects were recruited for 90 minute 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 3. Our aim was to obtain around 100 trading rounds per session, and as noted above, this was, in fact, the average number of trading rounds played in all sessions.

When we decided that the last game had been played, we asked our program to choose one game at random from all of the games played in that session. The subjects' final point totals from the chosen game were converted into a probability of winning a \$10 prize that was in addition to the \$10 payment that each subject was guaranteed to receive for participating in the 90–minute experimental session. Each additional point a subject earned above the initial 100 points they were given in each game increased their probability of winning the additional \$10 prize by the same amount if that game was the one chosen at random. The \$10 prizes were awarded based on a random choice for the cut–off probability. This mechanism for the awarding of additional \$10 prizes was carefully explained to subjects in the instructions and was intended to maximize their performance in every round of play. Subjects were further instructed that they were not competing with one another for the \$10 prizes. They were told that each player in the session could win a \$10 prize if he or she earned enough additional points in the game chosen, relative to the maximum number of points a player *of their type* could have been expected to earn in the game chosen. The maximum expected number of points takes into account a player's type, and the number of trading rounds played. Our use of a Roth–Malouf (1979) binary lottery to determine actual cash payments was intended to control for subjects' differing attitudes towards risk.

4 Experimental Results

We now turn to the results of our experimental sessions, focusing on the *aggregate* behavior by all players of a given type. Table 4 reports the aggregate frequency with which each type offers to

Table 4: Frequencies With Which Each Type Offers to Trade for the Good Corresponding to Their Type Over Each Half of a Session

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
A1	.99	.99	.97	.98	.94	.99
A2	.99	.96	.98	.92	.97	.99
A3	.95	.97	.97	.89	1.00	1.00
B1	.95	.87	1.00	.94	.97	.93
B2	.91	.97	1.00	.94	1.00	1.00
B3	.99	1.00	.93	.96	.90	1.00
C1	.97	.92	.92	.82	1.00	.94
C2	.95	.90	1.00	.99	1.00	1.00
C3	1.00	.99	1.00	.97	.98	1.00
D1	.95	.96	.93	1.00	.98	1.00
D2	.97	1.00	.99	.94	.99	1.00
D3	1.00	1.00	1.00	1.00	.98	1.00
D4	1.00	1.00	.98	1.00	.96	.91
D5	.94	.94	.74	.85	.74	.89
All Sessions[a]	.97	.96	.96	.94	.96	.97

[a]Weighted Average

trade for the good corresponding to their type over the first and second half of all 14 experimental sessions.[12] Since subjects were informed that they earned a positive net payoff in points only from acquiring the good corresponding to their type, the offer frequencies reported in Table 4 should all be 100%. Indeed, we see that in most cases, these frequencies are either at 100%, or very close to 100%. This finding serves to confirm that most subjects understood the role they were assigned to play as one of the three different types and were sufficiently well motivated by the incentive structure of our experimental design.

[12]If a session consisted of 100 trading rounds, the *first half* of the session consists of the first 50 trading rounds and the *second half* of the session consists of the remaining 50 trading rounds.

4.1 Observed behavior when good 0 has zero storage cost

Recall that when good 0 is least costly–to–store, there is one steady state in which good 0 is universally accepted in trade (by all three types) and another steady state, in which no one believes it has value, and therefore no type offers to trade for it. In either steady state, all three types are predicted to play "fundamental" strategies; speculative strategies should not be observed.

Tables 5abc report the aggregate frequencies with which each type *offers to trade for* good 0 over the first and second half of the 9 sessions in which good zero had no storage cost, and was therefore the least costly–to–store good (Sessions A1–A3, B1–B3, and C1–C3). The offer frequencies in these tables are further disaggregated according to the good that each type was offering in exchange for good 0.[13] As the tables reveal, when good 0 is least costly to store (has zero storage cost), it is widely, though not quite universally, accepted in exchange.

The offer frequencies in Tables 5abc, are in most cases far from zero, and by the second half of the sessions, are in many cases approaching 1.00, the prediction for the steady state in which good 0 serves as a medium of exchange. There are, however, a few exceptions. In particular, in sessions A3 and B2, type 1 players, become more reluctant to trade good 3 for good 0 from the first to the second half of these sessions. This behavior suggests that at least some type 1 players are "speculating" in good 3 (by refusing to trade good 3 for good 0), even though such speculative behavior never comprises a steady state best response in this parameterization of the model.[14] Notice also that by the second half of all 9 sessions, type 2 players are, in the aggregate, more willing to trade good 3 for good 0 than they are willing to trade good 1 for good 0; this difference in behavior is consistent with fundamental storage cost considerations, but in equilibrium, there would be no such difference.

The positive frequencies with which all types offer to trade for good 0, as reported in Tables 5abc, indicate that players believe good 0 has value. If good 0 is a medium of exchange in this environment, then players who hold it should only offer to trade it for their consumption good, and not for any other good. Indeed, as the left columns of Tables 6abc reveal, by the second half of each session, with few exceptions, players are rarely offering to trade good 0 for the good they

[13] We did not consider cases in which both players had good 0 in storage as there is no theoretical prediction in this case.

[14] Disaggregation of the data (from all sessions) reveals that most subjects are playing something close to a pure, rather than a mixed strategy when meeting a good that differs from the good they currently hold in storage. It is therefore appropriate to refer to the speculative behavior we observe as arising from the behavior of "some" type 1 players.

Table 5a: Frequencies With Which Type 1 Offers to Trade For Good 0, Sessions A1–C3.

		1 Offers to Trade for 0		1 Offers to Trade 2 for 0		1 Offers to Trade 3 for 0	
Session	$u =$	First Half	Second Half	First Half	Second Half	First Half	Second Half
A1	20	.89	.98	.90	.99	.80	.75*
A2	20	.63	.83	.67	.86	.43	.40
A3	20	.77	.83	.76	.86	.81	.56
B1	20	.99	1.00	.98	1.00	1.00	1.00
B2	20	.85	.89	.84	.97	1.00	.54
B3	20	.90	1.00	.93	1.00	.73	1.00*
C1	20	.84	.88	.85	.88	.60	†
C2	20	.83	.93	.85	.94	.69	.83
C3	20	.78	.92	.89	.92	.18	1.00*

Table 5b: Frequencies With Which Type 2 Offers to Trade For Good 0, Sessions A1–C3.

		2 Offers to Trade for 0		2 Offers to Trade 1 for 0		2 Offers to Trade 3 for 0	
Session	$u =$	First Half	Second Half	First Half	Second Half	First Half	Second Half
A1	20	.92	.97	.88	.90	.94	1.00
A2	20	.82	.80	.73	.73	.89	.84
A3	20	.87	.88	.68	.76	1.00	.96
B1	20	.95	.96	.91	.91	.97	1.00
B2	20	.93	.88	.87	.73	.97	.97
B3	20	.89	.91	.95	.90	.86	.91
C1	20	.76	.67	.64	.55	.79	.72
C2	20	.91	.96	.79	.93	.96	.98
C3	20	.88	1.00	.74	1.00	.96	1.00

Table 5c: Frequencies With Which Type 3 Offers to Trade For Good 0, Sessions A1–C3.

		3 Offers to Trade for 0		3 Offers to Trade 1 for 0		3 Offers to Trade 2 for 0	
Session	$u =$	First Half	Second Half	First Half	Second Half	First Half	Second Half
A1	20	.83	.95	.83	.95	.80	†
A2	20	.77	.82	.91	.88	.48	.60
A3	20	.53	.68	.57	.72	.37	.38
B1	20	.92	.85	.94	.87	.80	.50*
B2	20	1.00	1.00	1.00	1.00	1.00*	†
B3	20	.63	.70	.64	.71	.50	.67
C1	20	.92	.94	.93	.97	.90	.57
C2	20	.75	.59	.80	.68	.47	.06
C3	20	.65	.79	.67	.79	.40	.75*

* There were less than 5 observations for this trading situation.
† There were no observations for this trading situation.

neither consume nor produce.

Table 6a: Frequencies With Which Type 1 Offers to Trade For Good 3, Sessions A1–C3.

		1 Offers to Trade 0 for 3		1 Offers to Trade 2 for 3	
Session	$u =$	First Half	Second Half	First Half	Second Half
A1	20	.00	.05	.17	.11
A2	20	.06	.00	.30	.15
A3	20	.16	.10	.25	.13
B1	20	.00	.06	.24	.27
B2	20	.62	.38	.42	.48
B3	20	.06	.00	.32	.14
C1	20	.14	.00	.19	.00
C2	20	.10	.11	.40	.12
C3	20	.13	.00	.22	.21

Table 6b: Frequencies With Which Type 2 Offers to Trade For Good 1, Sessions A1–C3.

		2 Offers to Trade 0 for 1		2 Offers to Trade 3 for 1	
Session	$u =$	First Half	Second Half	First Half	Second Half
A1	20	.16	.03	.96	.95
A2	20	.04	.11	.82	.73
A3	20	.15	.07	.98	.94
B1	20	.11	.03	1.00	1.00
B2	20	.09	.08	1.00	1.00
B3	20	.03	.04	.90	.94
C1	20	.54	.19	.83	.85
C2	20	.26	.04	.84	1.00
C3	20	.11	.00	.96	1.00

Table 6c: Frequencies With Which Type 3 Offers to Trade For Good 2, Sessions A1–C3.

		3 Offers to Trade 0 for 2		3 Offers to Trade 1 for 2	
Session	$u =$	First Half	Second Half	First Half	Second Half
A1	20	.10	.00	.14	.06
A2	20	.21	.15	.33	.27
A3	20	.16	.09	.21	.16
B1	20	.19	.04	.19	.10
B2	20	.00	.00	.05	.00
B3	20	.13	.18	.28	.40
C1	20	.14	.06	.14	.10
C2	20	.12	.06	.43	.15
C3	20	.05	.08	.09	.05

The right–most columns of Tables 6abc reveal the frequencies with which each type is offering to trade their production good for the other, non–token, non–consumption good. Recall that, regardless of whether good 0 has value, the steady state prediction is that all three types adhere to "fundamental," storage–cost–reducing strategies in such trading encounters. For the most part, subjects appear to be adopting such strategies, especially by the second half of each session. In particular, type 1 and 3 players are refusing to trade their production goods 2 and 1 for the more costly–to–store goods 3 and 2, respectively, and type 2 players are offering to trade their production good 3 for the less costly–to–store good 1, which also serves as a medium of exchange in this environment. However, there are some exceptions to this finding; there is persistent speculation by some type 1 players in good 3 in sessions B1, B2, and C3, and persistent speculation by some type 3 players in good 2 in sessions A2 and B3. Similarly, in session A2, there is some deterioration in the frequency with which type 2 players trade good 3 for good 1.[15]

The only treatment variable that varied across sessions A1–A3, B1–B3 and C1–C3 was the value of m, the fraction of the population storing good 0. Recall that the probability that good 0 is accepted when offered is predicted to be invariant to changes in m, but that the analogous acceptance probability for good 1 is predicted to decrease as m increases. To assess whether these predictions hold, we calculated the number of times good 0 was offered in trade and the number of times that good 0 was both offered and accepted, i.e. the number of times good 0 was exchanged.[16] The ratio of the latter number to the former number is the frequency with which good 0 is accepted in trade when offered. We calculated this same acceptance frequency for good 1. These acceptance frequencies for the first and second halves of sessions A1–A3, B1–B3 and C1–C3 are presented in Table 7. We have also included in Table 7 the theoretical, equilibrium acceptance frequencies for goods 0 and 1 which are reproduced from Table 2 and are found under the heading Equil. Freq. In Table 7 we observe that, by the second half of each session, the frequency with which good 0 is

Table 7: Frequencies With Which Goods 0 and 1 are Accepted When Offered, Case 1 (Parameter Set 1), Sessions A1–C3.

		Good 0 Accepted When Offered			Good 1 Accepted When Offered		
Session	$m =$	First Half	Second Half	Equil. Freq.	First Half	Second Half	Equil. Freq.
A1	.25	.85	.95	1.00	.70	.63	.58
A2	.25	.74	.85	1.00	.73	.66	.58
A3	.25	.70	.82	1.00	.72	.66	.58
B1	.33	.97	.92	1.00	.68	.51	.52
B2	.33	.97	.90	1.00	.51	.64	.52
B3	.33	.76	.91	1.00	.64	.78	.52
C1	.50	.80	.81	1.00	.63	.52	.45
C2	.50	.81	.82	1.00	.62	.55	.45
C3	.50	.79	.86	1.00	.63	.57	.45

accepted when offered is always greater than 80%, but never exceeds 95%. These high, but less than

[15]Similar anomalous behavior was observed by Duffy and Ochs (1999) who examined this same experimental environment but without the token good 0. Again, disaggregation reveals that it is a certain fraction of players who are persistently playing pure, speculative strategies, and not all players adopting some kind of mixed strategy.

[16]We excluded from these calculations cases where both parties to an exchange had good 0 in storage.

universal acceptance frequencies are a reflection of the high offer frequencies reported in Tables 5abc. What is interesting to note is that, contrary to the theory, there appears to be some evidence that the acceptability of good 0 diminishes as m is increased. Using a simple nonparametric rank order test, we find no significant difference in acceptance frequencies for good 0 between the second half of sessions A1–A3 and B1–B3, and between the second half of sessions A1–A3 and C1–C3. However, we find that the acceptance frequencies in the second half of sessions C1–C3, where $m = .50$ are significantly lower ($p = .05$) than the corresponding frequencies for sessions B1–B3 where $m = .33$. Our data thus provide a mixed result with respect to the invariance of the acceptability of good 0 to changes in the value of m.

According to the theory, the acceptance of good 1 as a medium of exchange is strictly less than 100%, and the need for good 1 as a low cost medium of exchange decreases as m increases. The data in Table 7 indicate that by the second half of each session, the acceptability of good 1 is indeed less than that of good 0. However, in two out of three pair–wise comparisons over treatments, sessions A1-A3 versus B1–B3 and sessions B1–B3 versus C1–C3, robust rank order tests reveal that we cannot reject the null hypothesis that increases in m have no effect on the second half acceptance frequencies for good 1. Furthermore, in all but one session, (session B1) the second half acceptance frequencies for good 1 exceed the predicted equilibrium acceptance frequencies. These differences, however, might simply be attributed to the fact that the steady state is never perfectly achieved, as revealed in Tables 5 and 6.

In summary, when the token object has zero storage cost and is the least costly–to–store good, we observe that it is widely used as a medium of exchange, as predicted by the theory. However, there is a sufficient volume of unpredicted speculative behavior to call into question the sharper predictions of the theory. The token object is not universally accepted and it is not clear that increases in the supply of this object diminish agents' reliance on other goods as media of exchange.

4.2 Observed behavior when good 0 is costly–to–store

In experimental sessions D1–D5, we used parameter set 2 in which good 0 is more costly–to–store than good 1 but is much less costly–to–store than either good 2 or good 3.[17] The parameter values

[17] In sessions D1–D4, subjects were given the exact same information as in sessions A1–C3. In session D5, we provided subjects with additional information on the historical frequencies with which each type offered to trade good j for good $k \neq j$ in the current game. The provision of this additional strategic information appears to have had little effect on the aggregate outcomes as will become evident in Tables 7abc below.

chosen imply that in the steady state where good 0 is valued (which we focus on here) player types 1 and 2 will always trade their production goods 2 and 3, respectively, for either good 0 or good 1 so as to reduce their storage costs. Type 2 players who have traded for good 1 will be unwilling to trade it for good 0 and type 3 players, who produce good 1, will refuse to trade for any good other than good 3. Hence a type 1 player who holds good 0, can only obtain good 1 by first trading good 0 for good 3 (a speculative trade) and then trading good 3 to a type 3 player in exchange for good 1. In this version of the model, it is also a steady state best response for type 1 players to trade their production good 2 directly for good 3 (another speculative trade), thus bypassing the use of good 0 altogether. In either case, once a type 1 has traded for good 3, he should not offer to trade good 3 for good 0.

Tables 8abc display the relative frequencies with which each player type who faced the opportunity to trade for good 0 actually offered to make such a trade over the first and second half of sessions D1–D5. Consistent with the steady state prediction, both type 2 and type 3 players nearly always refuse to offer to trade good 1 for good 0, especially by the second half of a session. However, contrary to the steady state prediction, in the second half of four of the five sessions, 27–35% of all type 2 players who held good 3 and met good 0 in trade refused to offer to trade. As for the behavior of type 1 players, the data in Table 8a show that type 1 offers to trade good 2 for good 0 *decrease* somewhat from the first to the second half of four of the five sessions, and are never equal to the predicted frequency of 1.0. Also contrary to the steady state prediction, in the second half of three of the five sessions, one–third or more of type 1 players who had good 3 in storage and met good 0 in trade were willing to offer to trade their good 3 for good 0.

Finally, Tables 9abc report the frequencies with which each type offered to trade for the other, non–token, non–consumption good in sessions D1–D5. Recall that the steady state prediction is that type 1 will players will play the speculative strategy of offering to trade good 0 or good 2 for the more costly–to–store good 3, while types 2 and 3 continue to adhere to fundamental, storage cost reducing strategies. As shown in Table 9a, by the second half of each session, only 18 to 40% of type 1 players who held good 0 in storage and faced an opportunity to trade for good 3 offered to make the trade. Thus the majority of type 1 players refused to play the speculative strategy in this trading situation, despite the fact that by the second half of each session (as revealed earlier in Tables 8b and 8c), nearly all type 2 and 3 players refused to offer to trade good 1, the good desired by type 1 players, for good 0. By contrast, type 1 players did show a pronounced willingness

Table 8a: Frequencies With Which Type 1 Offers to Trade For Good 0, Sessions D1–D5.

		1 Offers to Trade 2 for 0		1 Offers to Trade 3 for 0	
Session	$u =$	First Half	Second Half	First Half	Second Half
D1	100	.92	.85	1.00*	.29
D2	100	.92	.79	.29	.33
D3	100	.92	.92	.67*	.80
D4	100	.72	.54	.67*	.40
D5	100	.84	.81	.63	.23

Table 8b: Frequencies With Which Type 2 Offers to Trade For Good 0, Sessions D1–D5.

		2 Offers to Trade 1 for 0		2 Offers to Trade 3 for 0	
Session	$u =$	First Half	Second Half	First Half	Second Half
D1	100	.09	.13	.90	.73
D2	100	.25	.00	1.00	1.00
D3	100	.24	.00	.78	.73
D4	100	.21	.00	.90	.70
D5	100	.07	.06	.73	.65

Table 8c: Frequencies With Which Type 3 Offers to Trade For Good 0, Sessions D1–D5.

		3 Offers to Trade 1 for 0		3 Offers to Trade 2 for 0	
Session	$u =$	First Half	Second Half	First Half	Second Half
D1	100	.14	.00	1.00*	†
D2	100	.02	.00	1.00*	†
D3	100	.02	.00	†	†
D4	100	.02	.00	.17	1.00*
D5	100	.13	.12	.40	†

* There were less than 5 observations for this trading situation.
† There were no observations for this trading situation.

to speculate by trading good 2 for good 3, a move that would increase their storage costs only slightly.[18] Thus, the aggregate pattern of behavior exhibited by the majority of type 1 players who faced an opportunity to trade for good 3 was to play fundamental except when they could get good 3 at little incremental cost.

Table 9a: Frequencies With Which Type 1 Offers to Trade For Good 3, Sessions D1–D5.

		1 Offers to Trade 0 for 3		1 Offers to Trade 2 for 3	
Session	$u =$	First Half	Second Half	First Half	Second Half
D1	100	.33	.40	.50	.70
D2	100	.35	.37	.57	.90
D3	100	.00	.18	.13	.30
D4	100	.10	.32	.55	.71
D5	100	.05	.33	.44	.55

Table 9b: Frequencies With Which Type 2 Offers to Trade For Good 1, Sessions D1–D5.

		2 Offers to Trade 0 for 1		2 Offers to Trade 3 for 1	
Session	$u =$	First Half	Second Half	First Half	Second Half
D1	100	.89	1.00	.88	.94
D2	100	1.00	1.00	1.00	1.00
D3	100	.95	1.00	.95	.93
D4	100	1.00	1.00	.88	1.00
D5	100	.92	1.00	.81	.84

Table 9a: Frequencies With Which Type 3 Offers to Trade For Good 2, Sessions D1–D5.

		3 Offers to Trade 0 for 2		3 Offers to Trade 1 for 2	
Session	$u =$	First Half	Second Half	First Half	Second Half
D1	100	.00	.00	.03	.00
D2	100	.08	.00	.05	.02
D3	100	.00	.00	.06	.00
D4	100	.33	.00	.10	.02
D5	100	.22	.14	.11	.09

Tables 9b and 9c reveal that types 2 and 3 were both adhering to the predicted fundamental strategies in the five D sessions. In particular, type 2 players were nearly always offering to trade goods 3 and 0 for the least costly–to–store good 1, and type 3 players were unwilling to trade goods 1 and 0 for the more costly–to–store good 2.

4.3 Were Payoff Incentives Sufficiently Salient?

One might question whether the incentives for type 1 players to speculate in good 3 under parameter set 2 were sufficiently large. One way to answer this question is to suppose that all three types were

[18]Recall that in parameter set 2, the difference in storage cost between goods 2 and 3 is just one point. ($c_2 = 23$ points, $c_3 = 24$ points). The 70% and higher rates of speculation that we observe in three of the five D sessions are the highest rates of speculation that have been reported in all of the experimental studies of the Kiyotaki–Wright (1989) model. This finding suggests that it is possible to obtain speculative behavior in the laboratory provided that storage cost differences are sufficiently small.

playing fundamental strategies in this environment. In that event, the steady state probability that each type successfully trades for his consumption good in any round is 1/9 (.111). However, if type 1 players play according to the "speculative" steady state equilibrium prediction, and speculate in good 3, while types 2 and 3 adhere to fundamental trading strategies, then the probability that each type successfully trades for his consumption good in any round is .157, an increase of 41%.[19]

A second way in which to address this question is to determine how many points subjects would have earned if they had played according to the predicted pure strategies and faced the same sequence of random matches they experienced in the laboratory sessions. This approach involves a simple simulation exercise, which we now describe. For each session A1–D5, we simulated the point totals that each subject would have earned at the end of each game had they played according to the steady state strategy profiles for the equilibrium in which good 0 serves as a medium of exchange (as discussed in section 2.2). In each simulation we used the same number of subjects, 18 or 24, that were present in a single experimental session. These simulated subjects were each endowed with the same initial good that the real subjects had in storage in the first round of every game. The simulated subjects also faced the same sequence of random matchings that were faced by their real counterparts. The length of each game was therefore the same as in the experimental session. We calculated the simulated point totals at the end of each game for each player, as it was end of game totals that determined players' earnings. We then calculated the ratio of the *actual* end-of-game point total earned by each real subject, to the simulated end-of-game point total for that player. Finally, we calculated the average ratio for all players of a given type over all games played in a session. These average ratios are reported in Table 10. Notice that in sessions A1-A3, B1-B3 and C1-C3, these average ratios are very close to 1.0, indicating that the real subjects' point totals were, on average, very close to what would have been predicted if they had always adhered to the fundamental trading strategies, and always offered to trade for good 0.

Note also that there is a substantial deterioration in these ratios for sessions D1-D4. The lower ratios in the D sessions are primarily a consequence of the lack of speculative play by Type 1 players

[19]Note further that if all three types play fundamental strategies in this case where good 1 is less costly to store than good 0, then one can show that in a steady state, type 1 players end up holding the entire stock of good 0! In this case, type 1 players are never able to trade good 0 for good 1 nor are they ever willing to trade good 0 for any other good. Since type 1 subjects earned a positive net payoff only when they successfully traded for good 1, it seems unlikely that they would have been content to hold onto good 0 for very long if they found they were unable to directly trade good 0 for good 1. For these reasons, we believe that there were strong incentives for type 1 players to adopt speculative strategies in all "D" sessions.

Table 10: Ratio of Subject Point Totals to Hypothetical Point Totals from Playing According to Steady State Pure Strategies: Averages for Each Player Type Calculated Using Point Totals at the End of All Games Played in a Session.

Session	Type 1	Type 2	Type 3
A1	0.95	1.00	0.99
A2	1.04	0.94	0.97
A3	0.83	0.97	0.93
All A	0.93	0.97	0.97
B1	0.95	0.97	0.97
B2	0.91	1.01	1.00
B3	1.02	0.97	0.99
All B	0.96	0.99	0.99
C1	0.93	0.95	0.95
C2	0.97	0.97	0.97
C3	0.97	0.97	0.98
All C	0.96	0.97	0.97
D1	0.89	0.81	0.91
D2	0.89	0.89	0.97
D3	0.59	0.63	0.89
D4	0.79	0.75	0.89
D5	0.64	0.63	0.77
All D	0.70	0.81	0.89

in these sessions. When Type 1 players played the speculative strategies of trading goods 0 and 2 for good 3 most frequently, as in session D2, the ratios in Table 10 are at their highest levels, and when Type 1 players played these speculative strategies least often, as in session D3, these ratios are at their lowest levels. The inefficient trading behavior of Type 1 players has spillover effects for Type 2 and 3 players as well, resulting in lower ratios for these player types.

Since the probability of winning the additional $10 prize was a linearly increasing function of the number of points earned in one randomly chosen round, the ratios in Table 10 also provide us with estimates of the average expected loss from not playing according to the predicted pure strategies in each session, (or over all sessions of a treatment). For instance, in all D sessions, Type 1 players earned on average, 70% of the points they could have earned had they played the steady state pure strategy profile. The average expected loss was therefore 30% of $10.00 or $3.00.

5 Conclusion

We have reported the results of an experiment that was designed to implement Kiyotaki and Wright's (1989) model of money as a medium of exchange with three goods and a single token object. The main question we sought to answer was whether an intrinsically worthless token object that was not invested with value by government decree, could nevertheless serve as a medium of exchange in accordance with the theory. The experimental evidence suggests that subjects will indeed offer to trade for a token object even though this object has no intrinsic value, is not legally required, and may not even be the least costly–to–store good. The equilibrium predictions of the Kiyotaki–Wright model come closest to being fulfilled when all types' best response is to play fundamental strategies, as in sessions A1–C3.

Nevertheless, even in these sessions, the frequencies with which goods 0 and 1 serve as media of exchange differ to some extent from the steady state equilibrium frequencies. In particular, Good 0 is accepted when offered *less* frequently and Good 1 is accepted when offered *more* frequently than predicted in the steady state. Furthermore, in contrast to the theory, the acceptability of the token object, good 0, appears to be sensitive to the supply of good 0, and the acceptability of the intrinsically valued good 1 may not decrease with increases in the supply of good 0. Finally, subjects have difficulty recognizing when speculative strategies are their best response in sessions D1–D5, despite adequate monetary incentives. In particular, many type 1 players who trade for good 0 fail to recognize that they will be unable to trade good 0 directly for the good they desire, i.e. good 1, the least costly–to–store good. The majority of type 1 subjects refuse to trade good 0 for the more costly–to–store good 3 and to get their consumption good 1, they must agree to such trades. We conclude that the token object –good 0– does serve as a medium of exchange in our implementation of the theoretical environment, but it appears to have a more limited role as a medium of exchange than theory predicts.

Appendix 1

In this appendix we define and demonstrate the existence of the stationary pure strategy equilibria discussed in the text using the parameter sets we considered in the experiment as reported in Table 1. The laboratory environment differs from Kiyotaki and Wright's (1989) theoretical environment in one important respect: instead of having a continuum of agents of each type we have a *finite* population of N agents with $N/3$ of each type, where $N = 24$ or 18. We must therefore demonstrate that the Nash equilibria in the theoretical environment with a continuum of agents are also Nash equilibria in the environment we study with a finite population of agents.

We begin with the definition of an equilibrium, which is similar to that found in Kiyotaki and Wright (1989) and Aiyagari and Wallace (1992). As described in the text, no good may be freely disposed of and all exchanges consist of one–for–one swaps of goods in inventory. When an agent successfully trades for his consumption good, he immediately consumes that good and produces a unit of his production good. Let τ index a particular player type, 1, 2 or 3, and let η represent the set of player types. Let $s^{jk}_{\tau,i}(t) \in [0,1]$ denote the strategy of a type τ player with index number $i = 1, 2, ..., N/3$, who is storing good j and is randomly paired with a player storing good $k \neq j$ in round t. Let $s_{\tau,i}(t)$ denote the vector of player i's strategies at t. In equilibrium, all i players of type τ will have the same strategy vector, denoted $s_\tau(t)$. Let the set of such strategy vectors be denoted by $s(t)$. Following Kiyotaki and Wright, we focus on pure strategy equilibria, where $s^{jk}_\tau(t) \in \{0,1\}$. If $s^{jk}_\tau = 1$, the type τ player offers to trade good j for good k; if $s^{jk}_\tau = 0$, the type τ player refuses to trade good j for good k. Let $p^j_\tau(t)$ denote the proportion of players of type τ who are storing good j at the start of round t, and let $p(t)$ denote the vector of all such proportions. Define the term

$$a^j_\tau(t) = p^j_\tau(t) - p^j_\tau(t)\sum_\eta \sum_{j\neq k} p^k_\eta(t) s^{jk}_\tau(t) s^{kj}_\eta(t) + \sum_\eta \sum_{j\neq k} p^k_\tau(t) p^j_\eta(t) s^{kj}_\tau(t) s^{jk}_\eta(t), \qquad (A.1)$$

which represents the fraction of type τ agents with good j immediately following trading at time t. The first term on the right hand side of this expression is the fraction of type τ agents with good j prior to trading, while the next two terms represent the outflow and inflow, respectively, of such agents immediately following trading. The sequence $\{p(t)\}$ evolves according to $p(t+1) = F[p(t), s(t)]$, where elements of the map F are defined as follows:

$$p^j_\tau(t+1) = \begin{cases} 0 & \text{if } j = \tau \\ a^j_\tau(t) + a^\tau_\tau(t) & \text{if } j = \tau + 1 \text{ modulo 3}, \\ a^j_\tau(t) & \text{if } j \neq \tau,\ \tau+1 \text{ modulo 3}. \end{cases} \qquad (A.2)$$

where $a^\tau_\tau(t)$ represents new production of good $\tau+1$ modulo 3. Type τ agents who successfully trade for good τ immediately produce a unit of their production good $\tau + 1$ modulo 3, hence $p^\tau_\tau(t) = 0$ $\forall t$. The initial conditions are that $\sum_\tau p^0_\tau(1) = m$, and $\sum_\tau p^{\tau+1}_\tau(1) = 1 - m$, where m represents the exogenously given fraction of the entire population that is endowed with good 0. Since good 0 is neither produced nor freely disposable, in addition to (A.1-A.2) we also require $\sum_\tau p^0_\tau(t) = m$ $\forall t$.

An *equilibrium* starting from $p(1)$ is a sequence $\{p(t+1), s(t)\}$ such that

1. $p(t+1) = F[p(t), s(t)]$,

2. $s_{\tau,i}(t)$ is individually optimal for any player i given $\{s_{\tau,-i}(t),\ s_\eta(t),\ p(t+1)\}$.

where $s_{\tau,-i}(t)$ is understood to be the time t strategy pursued by all other type τ players excluding player i. Part 2 of this definition differs from that of Kiyotaki and Wright who only required that $s_\tau(t)$ be optimal for any agent of type τ. The difference is that, in our framework, a single player i

is non–atomistic, and hence may affect the proportions $p(t)$ by his choice of strategy. We will focus on stationary pure strategy equilibria as in Kiyotaki and Wright.

To assess whether strategies are individually optimal, it will be useful to introduce some additional notation. Let $v^j_{\tau,i}$ denote the expected utility to individual i of type τ storing good j immediately following trading at time t. In particular, define

$$v^j_{\tau,i}(t) = \begin{cases} -c_j + \beta \sum_\eta \sum_k p_{\eta,k}(t+1)\left[s^{j,k}_{\tau,i}(t+1)s^{k,j}_\eta(t+1)v^k_{\tau,i}(t+1)\right. \\ \qquad \left. + \left(1 - s^{j,k}_{\tau,i}(t+1)s^{k,j}_\eta(t+1)\right)v^j_{\tau,i}(t+1)\right] & \text{if } j \neq \tau, \\ u + \beta \sum_\eta \sum_k p_{\eta,k}(t+1)\left[s^{\tau+1,k}_{\tau,i}(t+1)s^{k,\tau+1}_\eta(t+1)v^k_{\tau,i}(t+1)\right. \\ \qquad \left. + \left(1 - s^{\tau+1,k}_{\tau,i}(t+1)s^{k,\tau+1}_\eta(t+1)\right)v^{\tau+1}_{\tau,i}(t+1)\right] & \text{if } j = \tau. \end{cases} \tag{A.3}$$

Recall that in our environment, the utility from consumption u, is common across goods and types while storage costs differ across goods $j = 0, 1, 2, 3$, but not across types. Furthermore, there are no consumption/production decisions for agents to make, nor is there any free disposal. The individual optimizing conditions can be defined as:

$$s^{jk}_{\tau,i}(t) \in \begin{cases} \{0\} & \text{if } v^j_{\tau,i}(t) > v^k_{\tau,i}(t), \\ \{1\} & \text{if } v^j_{\tau,i}(t) < v^k_{\tau,i}(t), \\ [0,1] & \text{otherwise.} \end{cases}$$

Following Kiyotaki and Wright, we will not consider the last case, where individuals play mixed strategies, and will instead focus on steady state (time-invariant), pure strategy equilibria. Accordingly, we will drop all further references to time. Since players always trade for their consumption good $s^{j\tau}_{\tau,i}(t) = 1$ for all t, and strategies are assumed to be symmetric we can write the stationary pure strategy vector of a type τ player compactly as

$$s_\tau = \left(s^{\tau+1,0}_\tau, s^{k,0}_\tau, s^{\tau+1,k}_\tau\right),$$

where $\tau + 1$ is type τ's production good and $k \neq 0$ is the other intrinsically valued good that type τ neither produces nor consumes.

For our parameterization of case 1, where $0 = c_0 < c_1 < c_2 < c_3$, we claim that, as in Kiyotaki and Wright's (1989) model with a continuum of each agent type, there exist only two stationary pure strategy equilibria in our model with a finite population of each agent type. In one of these steady states, good 0 is used as a generally accepted medium of exchange, good 1 also serves as a medium of exchange and the strategy vectors are $s_1 = (1, 1, 0)$, $s_2 = (1, 1, 1)$, $s_3 = (1, 1, 0)$. In the other stationary equilibrium, good 0 is never accepted in trade, so that $s_1 = (0, 0, 0)$, $s_2 = (0, 0, 1)$, $s_3 = (0, 0, 0)$. In this steady state, good 1 continues to serve as a medium of exchange. For our parameterization of case 2, where $0 = c_1 < c_0 < c_2 < c_3$, we claim there exists a stationary pure strategy equilibrium where $s_1 = (1, 0, 1)$, $s_2 = (1, 0, 1)$, $s_3 = (0, 1, 0)$. In this equilibrium, type 1 players speculate in the most costly-to-store good 3, which serves as a limited medium of exchange along with goods 1 and 0.

Our strategy for demonstrating the existence of such equilibria is as follows. We first verify that our parameter sets 1 and 2 imply the existence of the stationary pure strategy equilibria referred to above, under the assumption of a continuum of agents of each type as in Kiyotaki and Wright (1989) and Aiyagari and Wallace (1992). We then turn our attention to the case of a finite population of N agents. We assume that $N - 1$ of these agents play according to the pure strategies that characterize one of the stationary equilibria of the model with continuum of agents. The lone remaining (N^{th}) agent plays some other strategy. We then numerically determine the expected population proportion of agents who are type $\tau = 1, 2, 3$ and storing good $j = 0, 1, 2, 3$, p^j_τ (a proxy for the steady state proportions). Denote these expected values by $\hat{p}^j_\tau$. Using these expected values, we then calculate the expected discounted utility to the deviant agent i of type τ from holding good j, $v^j_{\tau,i}$. We then rank the values $v^j_{\tau,i}$ and verify that this ranked list would imply that the deviant agent would do better by playing according to the pure strategy that is being played by the other type τ agents, which is the candidate stationary equilibrium strategy vector. We conduct this exercise for a single deviant agent of each of the three player types, for all values of N and m that we used in the experiment. This procedure allows us to confirm the existence of a stationary pure strategy Nash equilibrium in the finite population case.

Parameter Set #1

We first verified that our parameter set 1, where $0 = c_0 < c_1 < c_2 < c_3$, was consistent with the existence of a pure strategy equilibrium with a continuum of agents of each type, where all three types play fundamental strategies, and where good 0, the least costly–to–store good, serves as a generally accepted medium of exchange. Using the strategies $s_1 = (1, 1, 0)$, $s_2 = (1, 1, 1)$, and $s_3 = (1, 1, 0)$ in (A.1) we solved for the steady state proportions $p^j_\tau(t) = \bar{p}^j_\tau$ $\forall t$ using (A.2) for each of the three different values of m in parameter set 1. These steady state proportions are given in Table A1 below. Recall that p^j_τ and m represent proportions of the *entire population of players*, rather than proportions within a certain type (as in the original Kiyotaki and Wright (1989) paper). Using these steady state proportions together with the three pure strategy vectors and the model

Table A1: Steady State Inventory Proportions:
The Case Where Good 0 is Least Costly to Store, is Used as Money and
There is a Continuum of Each Player Type.

Case 1 (Parameter Set 1)

$m =$	$\bar{p}^0_1$	$\bar{p}^2_1$	$\bar{p}^3_1$	$\bar{p}^0_2$	$\bar{p}^1_2$	$\bar{p}^3_2$	$\bar{p}^0_3$	$\bar{p}^1_3$	$\bar{p}^2_3$
.25	.066	.267	.000	.082	.103	.148	.102	.231	.000
.33	.092	.241	.000	.109	.084	.140	.132	.201	.000
.50	.150	.183	.000	.163	.052	.118	.186	.147	.000

parameters given in parameter set 1, we used (A.3) to calculate the stationary (expected) value to each type of storing each good for each value of m. It is readily verified that in all cases, the rankings of these values satisfy:

$$\begin{array}{ccccccc} v^3_1 & < & v^2_1 & < & v^0_1 & < & v^1_1, \\ v^3_2 & < & v^1_2 & < & v^0_2 & < & v^2_2, \\ v^2_3 & < & v^1_3 & < & v^0_3 & < & v^3_3. \end{array}$$

These rankings imply that the strategy vectors $s_1 = (1, 1, 0)$, $s_2 = (1, 1, 1)$ and $s_3 = (1, 1, 0)$ are optimal for each type, given $\bar{p}$ and m.

Following the same procedure, we verified that the case where good 0 is never accepted in exchange and players adhere to the strategies $s_1 = (0, 0, 0)$, $s_2 = (0, 0, 1)$ and $s_3 = (0, 0, 0)$ is also a stationary pure strategy equilibrium under parameter set 1. Table A2 gives the steady state proportions in this case for the three different values of m found in parameter set 1. In our experimental environment, m was initially distributed evenly over the three types, hence $\bar{p}^0_\tau = m/3$ $\forall\tau$. Again, one can use these steady state proportions, the three strategy vectors, and the parameter

Table A2: Steady State Inventory Proportions:
Case Where Good 0 is Not Used as Money and There is a Continuum of Each Player Type.
Case 1 (Parameter Set 1)

$m =$	$\bar{p}_1^0$	$\bar{p}_1^2$	$\bar{p}_1^3$	$\bar{p}_2^0$	$\bar{p}_2^1$	$\bar{p}_2^3$	$\bar{p}_3^0$	$\bar{p}_3^1$	$\bar{p}_3^2$
.25	.083	.250	.000	.083	.125	.125	.083	.250	.000
.33	.111	.222	.000	.111	.111	.111	.111	.222	.000
.50	.167	.167	.000	.167	.083	.083	.167	.167	.000

choices in parameter set 1 in (A.3) to verify that for each value of m, $v_\tau^0 = 0 \; \forall \tau$, and:

$$\begin{aligned} v_1^3 &< v_1^2 < v_1^1, \\ v_2^3 &< v_2^1 < v_2^2, \\ v_3^2 &< v_3^1 < v_3^3. \end{aligned}$$

We next need to verify that in finite populations with 8 or 6 players of each type, where all players but one play according to these fundamental strategies, the expected distribution of goods over player types continues to imply the same ranking of the values v_τ^j. This last step cannot be done analytically (as far as we know). Consequently we resorted to numerical simulations. In particular, we considered populations of size $N = 24$ or 18 players (8 or 6 players of each type). We required $N-1$ players to made trading decisions according to the candidate fundamental equilibrium trading strategies. The remaining player played according to one of the eight possible pure strategy vectors for each player type. This set of eight strategy vectors includes the candidate equilibrium strategy as well as 7 other possible "deviant" strategies.

In the simulations, after the N players submitted offers, any mutually agreeable trades were implemented and the resulting changes in inventory distributions were recorded. In these simulations, as in the experiment, a fraction m of each player type initially has good 0 in storage; the remaining fraction of each player type initially has their production good in storage. Each simulation run consists of 10 rounds, the expected length of a game in our experiment. At the end of each run we calculated the average simulated proportions $\hat{p}_\tau^j$ over all 10 rounds played (so these values take into account our initialization procedure). We conducted a total of 10 simulation runs (as each experimental session consisted of approximately 10 games) for each value of m, population size, $N = 24$ or 18, and for the case where one player played according to one of the eight possible pure strategies for his type while the other $N-1$ players played according to the candidate equilibrium strategies. As an illustration, we report in Table A2, the mean and (standard deviation) of 10 $\hat{p}_\tau^j$ values for each value of m and N, for the case where a type 1 deviant – player number 1 – plays one of the 8 possible pure strategies s_{11}, and the other $N-1$ players adhere to the candidate equilibrium (fundamental) strategies.

Note that the first "deviant" strategy listed for each parameterization, $s_{11} = (1, 1, 0)$ is in fact, the candidate fundamental pure strategy; the proportions in this case can be compared with the theoretical steady state proportions given in Table A1. The difference between the proportions in Tables A3 and A1 are quite small and are due to the finiteness of the population. The other 7 strategies truly represent deviant strategies played by one type 1 deviant player. When these strategies are played, we also observe that there is not much change in the mean proportions relative to the theoretical steady state proportions.

Table A3: Mean $\hat{p}_\tau^j$ and (Standard Deviation) From 10 Simulation Runs
Case 1 (Parameter Set 1) Finite Population of 24 or 18 Players, $m = .25$, .33 or .50.

m =	N =	Deviant s_{11}	$\hat{p}_1^0$	$\hat{p}_1^2$	$\hat{p}_1^3$	$\hat{p}_2^0$	$\hat{p}_2^1$	$\hat{p}_2^3$	$\hat{p}_3^0$	$\hat{p}_3^1$	$\hat{p}_3^2$
.25	24	(1,1,0)	.070	.263	.000	.085	.097	.151	.095	.238	.000
			(.023)	(.023)	(.000)	(.016)	(.015)	(.016)	(.021)	(.021)	(.000)
		(0,0,0)	.065	.268	.000	.091	.100	.142	.094	.239	.000
			(.015)	(.015)	(.000)	(.029)	(.025)	(.017)	(.022)	(.022)	(.000)
		(0,0,1)	.055	.268	.010	.089	.100	.144	.106	.227	.000
			(.010)	(.013)	(.010)	(.019)	(.022)	(.021)	(.021)	(.021)	(.000)
		(0,1,0)	.058	.275	.000	.082	.108	.143	.110	.223	.000
			(.012)	(.012)	(.000)	(.028)	(.029)	(.017)	(.024)	(.024)	(.000)
		(0,1,1)	.071	.250	.012	.081	.103	.149	.097	.236	.000
			(.011)	(.016)	(.009)	(.019)	(.014)	(.020)	(.022)	(.022)	(.000)
		(1,0,0)	.058	.275	.000	.095	.088	.150	.098	.235	.000
			(.013)	(.013)	(.000)	(.016)	(.018)	(.013)	(.022)	(.022)	(.000)
		(1,0,1)	.055	.261	.017	.088	.115	.130	.107	.226	.000
			(.013)	(.017)	(.012)	(.016)	(.020)	(.014)	(.012)	(.012)	(.000)
		(1,1,1)	.061	.262	.010	.087	.106	.140	.102	.231	.000
			(.017)	(.019)	(.009)	(.018)	(.029)	(.033)	(.026)	(.026)	(.000)
.33	18	(1,1,0)	.092	.241	.000	.112	.073	.148	.129	.204	.000
			(.019)	(.019)	(.000)	(.034)	(.021)	(.026)	(.028)	(.028)	(.000)
		(0,0,0)	.085	.248	.000	.116	.072	.145	.132	.201	.000
			(.019)	(.019)	(.000)	(.034)	(.023)	(.025)	(.043)	(.043)	(.000)
		(0,0,1)	.082	.231	.020	.125	.069	.139	.126	.207	.000
			(.022)	(.027)	(.017)	(.023)	(.016)	(.026)	(.032)	(.032)	(.000)
		(0,1,0)	.075	.258	.000	.127	.077	.129	.131	.202	.000
			(.022)	(.022)	(.000)	(.022)	(.025)	(.031)	(.028)	(.028)	(.000)
		(0,1,1)	.072	.252	.009	.124	.088	.121	.137	.196	.000
			(.023)	(.024)	(.006)	(.024)	(.021)	(.022)	(.030)	(.030)	(.000)
		(1,0,0)	.094	.239	.000	.102	.088	.143	.137	.196	.000
			(.019)	(.019)	(.000)	(.014)	(.029)	(.025)	(.024)	(.024)	(.000)
		(1,0,1)	.108	.201	.023	.107	.080	.146	.118	.215	.000
			(.013)	(.023)	(.021)	(.032)	(.019)	(.028)	(.027)	(.027)	(.000)
		(1,1,1)	.098	.218	.017	.106	.081	.146	.129	.204	.000
			(.015)	(.018)	(.013)	(.025)	(.028)	(.025)	(.024)	(.024)	(.000)
.50	18	(1,1,0)	.147	.186	.000	.169	.041	.123	.184	.149	.000
			(.024)	(.024)	(.000)	(.022)	(.022)	(.020)	(.029)	(.029)	(.000)
		(0,0,0)	.139	.194	.000	.175	.043	.115	.186	.147	.000
			(.025)	(.025)	(.000)	(.027)	(.029)	(.029)	(.027)	(.027)	(.000)
		(0,0,1)	.130	.178	.025	.180	.040	.113	.190	.143	.000
			(.019)	(.029)	(.018)	(.038)	(.027)	(.031)	(.023)	(.023)	(.000)
		(0,1,0)	.136	.197	.000	.193	.048	.092	.171	.162	.000
			(.023)	(.023)	(.000)	(.019)	(.031)	(.022)	(.025)	(.025)	(.000)
		(0,1,1)	.146	.175	.012	.169	.053	.111	.185	.148	.000
			(.022)	(.029)	(.013)	(.023)	(.014)	(.020)	(.030)	(.030)	(.000)
		(1,0,0)	.146	.188	.000	.163	.040	.131	.191	.142	.000
			(.017)	(.017)	(.000)	(.021)	(.019)	(.019)	(.025)	(.025)	(.000)
		(1,0,1)	.153	.166	.014	.168	.051	.114	.179	.154	.000
			(.030)	(.018)	(.023)	(.029)	(.021)	(.033)	(.029)	(.029)	(.000)
		(1,1,1)	.154	.163	.016	.173	.044	.116	.173	.160	.000
			(.013)	(.019)	(.014)	(.029)	(.012)	(.020)	(.030)	(.030)	(.000)

For each deviant strategy and set of mean proportions reported in Table A2 we used expressions A1-A2 to calculate the value to the type 1 deviant of the four goods. In calculating these values, we assumed that the other $N-1$ players adhered to the candidate equilibrium strategies. It is simple but tedious to verify that in all cases the values to the type 1 deviant of the four goods always satisfy:

$$v_{11}^3 < v_{11}^2 < v_{11}^0 < v_{11}^1$$

indicating that the deviant's optimal strategy is to play the equilibrium strategy for type 1 players, $s_1 = (1,1,0)$.

Of course, we have also repeated this same exercise for the case of a single type 2 or type 3 deviant, though we do not report the results here. While these deviant players also affect the proportions of players storing the various goods, the optimal strategy for the type 2 and 3 players nevertheless remain the fundamental equilibrium strategies $s_2 = (1,1,1)$ and $s_3 = (1,1,0)$, as is revealed in an exercise similar to the one conducted above for the type 1 deviant player.

Finally we have repeated the same procedures described above for the other pure strategy equilibrium, where good 0 is never accepted in exchange, so that the strategy vectors are $s_1 = (0,0,0)$, $s_2 = (0,0,1)$ and $s_3 = (0,0,0)$. We found that in our populations of $N = 24 or 18 agents$, a single deviant player of each type, playing all possible strategies did not alter the expected inventory proportions in such a way as to reverse the rankings in the values to storing each good that would prevail in the stationary equilibrium where all players were playing according to the candidate equilibrium strategies. Hence, $s_1 = (0,0,0)$, $s_2 = (0,0,1)$ and $s_3 = (0,0,0)$ also comprise a stationary Nash equilibrium in our environment with a finite population of players.

Parameter Set #2

We next verify that our parameter set 2, where $0 = c_1 < c_0 < c_2 < c_3$, is consistent with the existence of a stationary pure strategy equilibrium where a continuum of agents of each type play according to the strategies $s_1 = (1,0,1)$, $s_2 = (1,0,1)$ and $s_3 = (0,1,0)$, and where good 0 serves a limited role as a medium of exchange even though it has higher storage cost than good 1. In this equilibrium, good 3 also has a limited role as a medium of exchange and good 1 serves as a generally accepted medium of exchange.

Following the same procedures as described above we used the three candidate strategy vectors and our choice of $m = 1/6$ in parameter set 2 to calculate the steady state inventory proportions which are given in Table A4. Using these steady state proportions, the three strategy vectors and

Table A4: Steady State Inventory Proportions:
The Case Where Good 0 is Not Least Costly to Store but is Used as Money
and There is a Continuum of Each Player Type.

Case 2 (Parameter Set 2)

$m =$	$\bar{p}_1^0$	$\bar{p}_1^2$	$\bar{p}_1^3$	$\bar{p}_2^0$	$\bar{p}_2^1$	$\bar{p}_2^3$	$\bar{p}_3^0$	$\bar{p}_3^1$	$\bar{p}_3^2$
.167	.107	.157	.069	.060	.186	.087	.000	.333	.000

the parameterization of the model given in parameter set 2, one can use the definitions in (A.3) to verify that:

$$\begin{array}{ccccccc} v_1^2 & < & v_1^0 & < & v_1^3 & < & v_1^1, \\ v_2^3 & < & v_2^0 & < & v_2^1 & < & v_2^2, \\ v_3^2 & < & v_3^0 & < & v_3^1 & < & v_3^3. \end{array}$$

These rankings imply that the strategy vectors $s_1 = (1,0,1)$, $s_2 = (1,0,1)$ and $s_3 = (0,1,0)$ would be optimal for each type, given $\bar{p}$ and m.

We then repeated the exercise we conducted for case 1, to verify that this stationary equilibrium remains an equilibrium with a finite population of $N = 18$ agents (6 of each type). In particular, we considered the possibility that one agent of each type played each of the eight possible strategies for his type, while the other $N-1$ agents adhered to the candidate equilibrium strategies. We calculated the expected inventory proportions and used these to determine the rankings of the value functions for the deviant player. In all cases we found, once again, that the rankings of these value functions implied that the deviant player would find it optimal to adhere to the candidate equilibrium strategy for his type. We therefore conclude that the strategy vector $s_1 = (1,0,1)$, $s_2 = (1,0,1)$ and $s_3 = (0,1,0)$ constitutes a stationary Nash equilibrium in our environment with a population of just $N = 18$ players. We note that, analogous to case 1, there is another stationary Nash equilibrium under parameter set 2 where the token object is not accepted in exchange, and the strategy vectors are $s_1 = (0,0,1)$, $s_2 = (0,0,1)$ and $s_3 = (0,0,0)$. We have also verified the existence of this equilibrium with a finite population of agents following the same procedures as outlined above for the other cases.

Appendix 2

This appendix provides the written instructions that were given to subjects in experimental sessions D1–D4. 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 experiment 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 four different types of goods: good 0, good 1, good 2, and good 3.

Each player begins each game with 100 points. In addition, at the start of a game each player has either one unit of good 0 or 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

In this session, 1 player of each type starts a game with good 0 in storage. The remaining players start a game with the good they "produce" in storage.

Players may continue to hold good 0 or the good they produce in storage, or they may exchange these goods for one unit of the good held in storage by another player. Exchanges of goods 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 a many units as possible of good 2, and type 3 players will want to obtain as many units as possible of good 3. Storing certain goods will cost you points. Storing good 0 costs you 1 point per round, storing good 2 costs you 23 points per round and storing good 3 costs you 24 points per round.

There is no cost for storing good 1. 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	-	23	=	77 points
2	100	-	24	=	76 points
3	100	-	0	=	100 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 at the top of the next page. The good you now have in storage is the good you received from the other player.

Table 2

Good	Storage Cost Per Round
0	1 point
1	0 points
2	23 points
3	24 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 either good zero or the good your type produces 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 four 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 90 minute 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.R. and N. Wallace (1991), "Existence of steady states with positive consumption in the Kiyotaki–Wright model," *Review of Economic Studies* 58, 901–916.

Aiyagari, S.R. and N. Wallace (1992), "Fiat money in the Kiyotaki–Wright model," *Economic Theory* 2, 447–464.

Bernasconi, M. and O. Kirchkamp (2000), "Why do monetary policies matter? An experimental study of saving and inflation in an overlapping generations model," *Journal of Monetary Economics* 46, 315–343.

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

Duffy, J. (1998), "Monetary theory in the laboratory," *Federal Reserve Bank of St. Louis Review* 80 (September/October), 9–26.

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

Holland, J.H. (1986), "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, 593–623.

Kiyotaki, N. and R. Wright (1989), "On money as a medium of exchange," *Journal of Political Economy* 97, 927–54.

Lian, P. and C.R. Plott (1998), "General equilibrium, markets, macroeconomics and money in a laboratory experimental environment," *Economic Theory* 12, 21–75.

Lim, S.S., E.C. Prescott, and S. Sunder (1994), "Stationary solution to the overlapping generations model of fiat money: experimental evidence," *Empirical Economics* 19, 255–277.

Marimon, R., E. McGrattan and T.J. Sargent (1990), "Money as a medium of exchange in an economy with artificially intelligent agents," *Journal of Economic Dynamics and Control* 14, 329–73.

Marimon, R. and S. Sunder (1993), "Indeterminacy of equilibria in a hyperinflationary world: experimental evidence," *Econometrica* 61, 1073–1107.

Marimon, R. and S. Sunder (1994), "Expectations and learning under alternative monetary regimes: an experimental study," *Economic Theory* 4, 131–162.

McCabe, K.A. (1989), "Fiat money as a store of value in an experimental market," *Journal of Economic Behavior and Organization* 12, 215–231.

Ostroy, J.M. and R.M. Starr (1990), "The transactions role of money," Chapter 1 of B.M. Friedman and F.H. Hahn (eds.), *Handbook of Monetary Economics* vol. 1, New York: North–Holland, 3–62.

Roth, A.E. and M.W.K. Malouf (1979), "Game theoretic models and the role of information in bargaining: An experimental study," *Psychological Review* 86, 574–94.

Shi, S. (1995), "Money and prices: a model of search and bargaining," *Journal of Economic Theory* 67, 467-496.

Trejos, A. and R. Wright (1995), "Search, bargaining, money and prices," *Journal of Political Economy* 103, 118–141.

Figure 1: Illustration of the Trading Screens Used in the Experiment

A. Trading Round Screen # 1

You are a type **2** player. This is Round **5**. Your Point Total: **151.**

You earn 100 points per unit of good 2 you obtain.
It costs 1 point per round for storing good 0.
It costs 0 points per round for storing good 1.
It costs 23 points per round for storing good 2.
It costs 24 points per round for storing good 3.

You currently have good 0 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 0 for the other player's good 2? Y/N:

Probability the game will end:

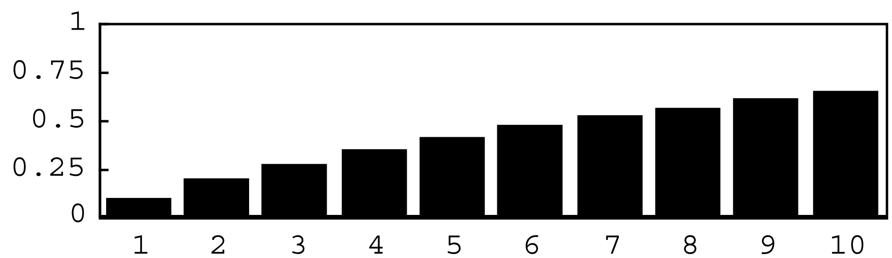

Rounds from now.

Percent of Each Type of Player Storing Each Type of Good
Historical Average as of Round 4

Type	Good 0	Good 1	Good 2	Good 3
1	0.27	0.00	0.67	0.06
2	0.20	0.33	0.00	0.47
3	0.03	0.97	0.00	0.00

B. Trading Round Screen # 2

You are a type **2** player. This is Round **5**. Your Point Total: **227.**

You earn 100 points per unit of good 2 you obtain.
It costs 1 point per round for storing good 0.
It costs 0 points per round for storing good 1.
It costs 23 points per round for storing good 2.
It costs 24 points per round for storing good 3.

In the last round you had good 0 in storage.
You were matched with a player of type 1.
This player had good 2 in storage.
You proposed to trade your good 0 for the other player's good 2.
The other player agreed to trade.

You received a net payoff of 76 points for the round.
Your new point total is 227.
You now have good 3 in storage.

Percent of Each Type of Player Storing Each Type of Good
Historical Average as of Round 5

Type	Good 0	Good 1	Good 2	Good 3
1	0.28	0.00	0.64	0.08
2	0.19	0.36	0.00	0.45
3	0.03	0.97	0.00	0.00

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