

John Duffy is an assistant professor of economics at the University of Pittsburgh and a visiting scholar at the Federal Reserve Bank of St. Louis. The author thanks Jack Ochs, Michael Pakko, and Joseph Ritter for helpful comments, Nick Meggos for research assistance, and Denise Hazlett and Shyam Sunder for permission to reproduce figures from their papers.



Monetary Theory in the Laboratory

John Duffy

INTRODUCTION

During the past few decades, an increasing number of economists have begun conducting controlled laboratory experiments with paid human subjects in an effort to test theories of economic behavior. Economists have turned to laboratory experiments because in many instances they lack the appropriate nonexperimental field data that are required to test their theories. Economic experiments usually involve groups of subjects, typically college students, who play simple economic decision-making games and are paid on the basis of the decisions they make. The games subjects play and the monetary rewards they can earn in the experiment are constructed in a manner that approximates the theoretical environment being studied so that the data collected from the experiment can be used to prove or disprove the theory.

To date, the experimental methodology has been used primarily by microeconomists and game theorists to test theories of individual or strategic behavior in various economic environments.¹ Using laboratory experiments, researchers have made substantial inroads in these fields; microeconomists and game theorists now readily acknowledge the relevance of using experiments to test theory and have begun to develop new theories to explain experimental findings that are at odds with existing theories.

By contrast, there has been relatively little use of the experimental methodology in macroeconomics. A likely explanation is

that the scale of macroeconomic phenomena is perceived to be too great to be tested in controlled laboratory experiments with small groups of subjects. That would almost certainly be the view, for example, among applied macroeconomists working with field data. Today, however, most macroeconomic theorists work with models that build on explicit microfoundations for the behavior of individuals. These models often impose artificial restrictions on individuals' behavior that are not embodied in any available field data. Consequently, it often is not possible to empirically test the aggregate predictions that emerge from these models. Even in those cases where the aggregate predictions of microfoundation models can be tested using field data, it is not always possible to use field data to verify whether behavior at the individual level adheres to the predictions or assumptions of these models.

On the other hand, controlled laboratory experiments can be used to test both the aggregate and individual predictions of microfoundation models. Though economists need to be wary of generalizing from the results of laboratory experiments that involve small groups of subjects to aggregate macroeconomic activity, there should be somewhat less objection to using experiments to test the predictions of individual behavior in macroeconomic models with explicit microfoundations. Because the predictions of these macroeconomic models stand on their assumptions of how individuals behave, laboratory evidence can test the reasonableness of a particular model as a characterization of some macroeconomic activity. And in the increasing number of macroeconomic models that admit multiple equilibria, laboratory evidence can be used to guide researchers toward those equilibria that subjects perceive as more relevant. Finally, macroeconomists can use laboratory methods to conduct policy experiments that would be impossible to conduct outside of the laboratory.

¹ For a recent survey, see, for example, Kagel and Roth (1995).

For these and other reasons, a few researchers have begun to use laboratory experiments, as well as survey questionnaires, to test both the assumptions and predictions of modern macroeconomic models. This paper describes some of these studies, focusing on laboratory tests of various different models of money.² Though the major emphasis is on the results obtained from laboratory studies of monetary theories, some attention is also devoted to issues of experimental design.

THE ROLES AND TYPES OF MONEY

Monetary theorists have no single unified theory of money and do not appear close to achieving any kind of consensus on what a reasonable theory might entail. Instead, there are several different and competing models of money. Each model can account for some, but not all, of the important roles that money plays in the economy. In reviewing laboratory tests of these theories it is useful to organize the discussion around the three major roles of money. Money serves as a *store of value*, as a *medium of exchange*, and as a *unit of account*. Money is a store of value because it is believed to retain purchasing power over time. Money is a medium of exchange because it is readily accepted in exchange for goods or services. Finally, money is a unit of account in that prices of goods and services are quoted all in terms of units of money.

In the discussion that follows, we will also make reference to two different types of money: *fiat money* and *commodity money*. Fiat money is unbacked, intrinsically worthless liabilities issued by the government and declared money by the government's fiat power. Dollar bills are an example of fiat money. Commodity money is an intrinsically valued good, the value of which serves as the value of money. Gold coins are an example of commodity money. Most of the experimental studies of money have focused on fiat money because that is the type of money most commonly used in the world today.

Money as a Store of Value

For money to serve as a store of value, economic agents must believe that money will have value in the future. A necessary condition for such a belief is that there be a future. Suppose it was known that the world would end at a certain future date. Would individuals continue to hold money? According to economic theory, no rational individual would choose to hold money the day before the world ended, as money would be worthless the next day. By the same kind of logic, no one would choose to hold money two days before the world ended, and therefore no one would choose to hold money three days before the world ended, and so on. This argument, based on *backward induction*, implies that if the end of the world were known with certainty, individuals would cease holding money. Do human subjects behave in accordance with this prediction? This is a good example of a question that is difficult to address outside of a controlled laboratory setting.

Testing Backward Induction

McCabe (1989) conducted an experiment that attempted to test for backward induction.³ In his experiment, a population of players is divided up equally into three types. The players then engage in a simple trading game that lasts six rounds. In each round of this game, two of the three player types are endowed with a unit of a good that yields a cash payoff at the end of the round, but which ceases to have any value in the following round. One can think of these goods as perishable, like fish or bananas. Furthermore, one out of every three players also possesses a ticket that can be used to buy one unit of any perishable good. Tickets are not perishable and can be potentially used over and over again to purchase perishable goods. Like fiat money, these tickets have no intrinsic value—they yield no cash payoff to their holders.⁴ In every round of the game, a ticket holder may use a ticket to buy one unit of any good, but only if there is another player willing to sell a good for a ticket. Barter exchanges of goods for goods are not

² Some of this literature is also discussed in surveys by Marimon (1997), Ochs (1995), and Sargent (1993).

³ While there have been many experiments involving tests of backward induction, (see, for example, Smith, Suchanek, and Williams [1988] and McKelvey and Palfrey [1992]), McCabe's (1989) is the only one to consider whether subjects apply backward induction in assessing the value of money.

⁴ Wallace (1998) proposes that money should not be a primitive in theories of money. In experimental tests of monetary theories there is a similar dictum: The goods that are candidates to play the roles of money should not be referred to as money; this prevents triggering any preconceived notions that experimental subjects might have regarding the proper role of a money object. Hence McCabe and other experimenters testing theories of money use more neutral terms like tickets to refer to objects that may serve as money.

allowed. All offers and exchanges are made anonymously through a centralized clearinghouse (the experimenter). After any exchanges occur, players are paid the cash value of any goods they are holding at the end of the round. The cash value from holding goods varies in a systematic way over the six rounds of the game and always differs across the three player types. Holding goods yields a cash payoff to each player type of either \$0.50, \$0.25, or \$0 depending on the round number, and this information is common knowledge. Similarly, the endowment pattern of goods varies in a systematic manner that is also common knowledge. The cash values of goods and the timing of the endowment of these goods are such that, if the game continued without end, all players would maximize their earnings by offering to buy or sell tickets in certain rounds of the game. That is, the use of tickets (fiat money) in this game is *Pareto improving*. The game, however, does not continue forever. In round 6, which all players know is the final round of the game, no player will want to remain holding a ticket because a ticket has no future use or cash value. By backward induction, no player should offer to trade for a ticket in any round and therefore no exchanges should ever be observed.

What happens in the experiment? McCabe reports that in contrast to the theoretical prediction of no exchange, there is a substantial volume of trade in tickets, however, this trading volume does diminish somewhat as the end of a six-round game approaches. When the same six-round game is repeated 10 to 20 times with the same group of subjects, these results are unchanged. It was only after McCabe brought the same group of subjects back for two additional experimental sessions, each a week apart, that he was able to obtain a collapse of all exchange in tickets from the start of a six-round game. Before McCabe obtained this result, subjects played the six-round game 45 times!

McCabe's experiment suggests that a fiat object, a ticket, may serve as a store of value for some period of time even when there is a certain, known finite horizon, if

players do not have sufficient experience with the implications that a finite horizon has for the value of their fiat money holdings. Though subjects clearly failed to apply backward induction, there are certain features of McCabe's experimental environment that may have served to promote the use of tickets as a store of value. In particular, the price of tickets in terms of goods was known to be fixed at one-for-one, so even though there was uncertainty regarding the acceptability of tickets, there was no uncertainty regarding the exchange rate of tickets for goods if tickets were accepted in trade. Moreover, tickets were the only means of transferring goods; barter exchanges of goods for goods were not allowed.

Overlapping Generations Models of Money

Historically, when fiat money has ceased to serve as a store of value, the explanation has not been that the fiat object was known to have zero value at a certain future date. Rather, the explanation is always that there has been a rapid erosion of money's purchasing power caused by inflation, as in a very high inflation, or *hyperinflation*. One theory of hyperinflation suggests that such events may be the result of *self-fulfilling expectations*. Individuals expect that future prices will be higher and make consumption and savings decisions in accordance with this expectation. Consequently, their expectations of future inflation become ratified in an actual inflation that is consistent with their expectations. Laboratory tests of this theory have been conducted using Samuelson's *overlapping generations* model of money, which we briefly summarize.

At every date t , a generation of n identical agents that live for two periods is born; in period t they are young and in period $t + 1$ they are old. Each agent consumes c_t^y units of a single, perishable consumption good when young and c_{t+1}^o units of this good when old. All agents are endowed with e_y units of the consumption good when young and e_o units of this good when old, with $e_y > e_o$. In addition, each member of the old generation has a constant m units of a durable but intrinsically worthless fiat object, or

money, that they acquired in their youth. These old agents may use their fiat money holdings to purchase some of the young agents' endowment of the consumption good at an endogenously determined per unit price of p_t . Because old agents will not be alive in the next period, they will trade all of their fiat money holdings for m/p_t units of the consumption good held by young agents.

Assuming that lifetime utility is given by $u(c_t^y, c_{t+1}^o) = c_t^y \cdot c_{t+1}^o$, and that agents have perfect foresight knowledge (rational expectations) of future prices, the representative young agent's problem is to

$$\max_{c_t^y, c_{t+1}^o} c_t^y \cdot c_{t+1}^o,$$

subject to a pair of budget constraints, one for each period of life as follows:

$$c_t^y = e_y - m_t / p_t,$$

$$c_t^o = e_o + m_t / p_{t+1}.$$

The first-order conditions imply that in equilibrium,

$$(1) \quad p_{t+1} = f(p_t) = -\frac{2m_t}{e_o} + \frac{e_y}{e_o} p_t.$$

Equation 1 is a first-order difference equation for the evolution of equilibrium price sequences. The specification of the model is completed by assuming (for now) that m_t is set equal to a constant $m > 0$ for all t . Given this assumption, the model has a unique stationary solution that is given as follows:


$$\bar{p} = \frac{2m}{e_y - e_o}.$$

An illustration of Equation 1 and the stationary solution is provided in Figure 1.

Figure 1 plots the price level at time $t + 1$ against the price level at time t . We see that unless p_t is equal to the stationary solution, \bar{p} , the equilibrium path of prices will be nonstationary. For instance, if $p_t = p' > \bar{p}$, the equilibrium path for prices, as illustrated in Figure 1, is explosive, and in the limit, fiat money holdings are worthless and no exchange occurs between young and old agents—they

Figure 1

A Self-fulfilling, Hyperinflationary Path for Prices



simply consume their endowments. In this case, individuals rationally forecast that the price level will rise by an ever-increasing amount in every period, and their forecasts are always correct—hence the name *self-fulfilling model of hyperinflation*.

Lucas (1986) noted that the multiplicity of nonstationary equilibrium paths for prices predicted by this model might be ruled out by assuming that agents did not possess rational expectations, but instead formed their expectations of future prices in an adaptive manner. In particular, if young agents' time t expectation of next period's price level, p_{t+1}^e , is given by some weighted average of past prices and forecasts, for example, if

$$p_{t+1}^e = \alpha p_{t-1} + (1 - \alpha) p_t^e, \quad \alpha \in (0, 1),$$

then one can show that the stationary solution, \bar{p} , will be the long-run price level of the economy. Replacing the assumption of rational expectations with adaptive expectations affects the stability properties of the stationary solution, changing it from an unstable equilibrium to a stable one.

Lucas argued that the question of which type of equilibrium path for prices was more likely (stationary or nonstationary-hyperinflationary) could be addressed in a controlled laboratory experiment. Following up on Lucas' suggestion,

Lim, Prescott, and Sunder (1994) designed an experiment that aimed to effectively implement the overlapping generations environment and address the question posed by Lucas. In every period of Lim et al.'s experimental sessions, a group of subjects of size n was chosen from an N -member subject pool to take a two-period turn in the model environment. Subjects entered first as a member of the young generation, and in the next period they were members of the old generation. Thus in every period of a session there were $2n$ subjects participating in the model environment and $N - 2n > n$ subjects who were idle. Each agent in the young generation was endowed with e_y chips and each old agent was endowed with e_o chips ($e_y > e_o$). In addition, the initial old generation of subjects (who, unlike other participants, were only in the model environment for one period), were each endowed with m francs. At the end of a two-period turn in the model environment, each subject earned cash payments that were proportional to the lifetime utility objective of $c^y \cdot c^o$, where c^y and c^o were the amount of chips a subject acquired when young and old, respectively. Subjects who were idle in any period played a forecasting game in which they earned cash payments according to how well they forecast the market clearing price level that would prevail in the next period, that is, how close their expectations of next period's price level, p_{t+1}^e , were to the actual price level p_{t+1} . Because of the backward induction problem, the experimenters did not want to reveal to subjects when the experimental session would end.⁵ When experimenters finally did announce the end of the session, they used the average forecast of the future price level, as determined by the nonparticipating subjects, to determine the value of all money holdings in terms of chips. They thus prevented the backward induction problem and had a final period price level that was endogenously determined by the subjects themselves.

After some experimentation, Lim et al. settled on a mechanism for the exchange of chips for francs between young and old

agents. Young agents were asked to provide a schedule of the quantity of chips they would supply at various different prices. The chip supply schedules from all young agents were then aggregated to form a market supply schedule. The market demand for chips was the amount of chips that could be bought with the total quantity of francs held by all members of the old generation (that is, old agents passively offered all of their francs for chips). The price at which aggregate supply equaled aggregate demand was announced as the market-clearing, equilibrium price level. This price was used to determine the quantity of chips supplied by each young agent (according to their individual supply schedules), the lifetime utility of old agents who would then exit the model environment, and the winner of the price forecasting game among those agents who were not participating as members of either the young or old generation.

In all of Lim et al.'s experimental sessions the price level fluctuated somewhat as agents learned how to make good decisions, but it was never explosive. Indeed, near the end of the experiments, the price level was close to, though never precisely equal to, the unique stationary value \bar{p} .⁶ The authors conclude from their experimental findings that stationary solutions for prices are good candidates for the equilibria of overlapping generations models and that explosive self-fulfilling hyperinflationary equilibria appear unlikely.

In some respects, this conclusion is not too surprising. As previously noted, Lucas (as well as others) argued that if individuals did not possess rational expectations of future prices and instead formed their expectations adaptively, then the stability of stationary equilibria such as \bar{p} might well be altered. Surprisingly, Lim et al. did not analyze or report forecasted prices from their experiment, preferring to concentrate their attention on the behavior of the market price level. They did note, however, that forecasts appeared to be adaptive, an observation that might account for their findings.

⁵ Of course, subjects were aware that the experimental session would last for a finite period—no more than a couple of hours. They did not know, however, precisely when the game they were playing would end.

⁶ In fact, the price level was even closer to another stationary Nash equilibrium where players take account of the size of the group of buyers and sellers and do not behave as if they are atomistic in their ability to affect prices. This equilibrium is close to the stationary equilibrium \bar{p} .

Setting aside the issue of whether individuals have rational expectations, one may question the self-fulfilling theory of hyperinflation that Lim et al. tested on other grounds. In particular, this theory ignores any causal relationship between the growth rate of the money supply and the rate of inflation of the price level. Inflation occurs even though the stock of money is constant. Historically, most hyperinflations have been ended by some effort to reduce the growth rate of the money stock (see, for example, Sargent [1986]), which suggests that a growing money stock should be a prerequisite for a hyperinflation to occur.

Marimon and Sunder (1993, 1994, and 1995) have conducted a series of experiments involving versions of the overlapping generations model where the stock of fiat money is allowed to grow over time. In these experiments there exists a government (the experimenter) that seeks to purchase a constant per capita amount, d , of the consumption good (chips) in every period. The government purchases d chips at price p_t by expanding the money supply (the amount of francs per capita) m as follows:

$$(2) \quad \Delta m = m_t - m_{t-1} = p_t d.$$

Using the policy rule, Equation 2, in combination with Equation 1 and assuming rational expectations, one can derive a law of motion for *gross inflation*, defined by $\pi_t = p_t / p_{t-1}$, of the form:

$$\pi_{t+1} = g(\pi_t).$$

This equation is illustrated in Figure 2.⁷

In this version of the model, if d is small enough, there exist two stationary inflation rates, $\pi_{low} < \pi_{high}$. Under the assumption that agents possess rational expectations, the higher of these two stationary values is predicted to be the limit of all equilibrium sequences for inflation starting from any initial value for $\pi_t > \pi_{low}$; Figure 2, illustrates one such equilibrium sequence with an initial inflation rate of π' .

What happens when subjects are placed in this environment? Do they coordinate their expectations in such a way as to achieve π_{high} , π_{low} or some explosive path for inflation? Marimon and Sunder (1993) sought to address these questions using an experimental design that was similar to that of Lim et al. (1994). One difference was that nonparticipating subjects now forecast the future *inflation rate* rather than the price level. Furthermore, the government's means of financing its consumption requirement by issuing more francs in every period was made common knowledge.

Marimon and Sunder report that in contrast to the rational expectations prediction, the actual path for inflation in all of their experimental sessions is always in the neighborhood around the low-inflation stationary equilibrium, π_{low} . In particular, inflation rates were never observed around the high-inflation steady state, π_{high} , nor were they ever explosive. Figure 3 provides a representative illustration of the actual inflation path from one of Marimon and Sunder's experimental sessions (No. 7C). Marimon and Sunder's findings using the policy rule, Equation 2, have been replicated in an independent experimental study conducted by Arifovic (1995).⁸

As in Lim et al.'s experiment, Marimon and Sunder's (1993) finding can be explained by subjects' use of some kind of adaptive rather than rational expectations.⁹ Marimon and Sunder explore this possibility and conclude that though no single adaptive rule explains behavior in all of their experiments, a rule for expected inflation, π^e , of the type suggested by Cagan (1956), that is, $\pi_{t+1}^e = \pi_t^e + \alpha(\pi_{t-1} - \pi_t^e)$, $0 < \alpha < 1$, fits the data from many sessions reasonably well. When estimated versions of such rules are used in place of rational expectations, the low-inflation steady state, π_{low} , is locally stable and the high inflation steady state, π_{high} , is locally unstable, thereby explaining the attractiveness of π_{low} .¹⁰

Using the same experimental design, Marimon and Sunder (1994) conducted further overlapping generations experiments in an effort to understand the effect that

⁷ The derivation of this equation makes use of the fact that

$$m_t / p_t = e_y - \hat{c}_t,$$

where \hat{c}_t is the young agent's optimal consumption decision. The function

$$g(\pi_t) = \frac{e_o + e_y - 2d}{e_o} - \frac{e_y}{e_o \pi_t}.$$

⁸ Experimental results are replicated when qualitatively similar results are obtained using the same experimental design and treatment variables but with different groups of subjects. This notion of replications differs from that of Dewald, Thursby and Anderson (1986), where replication refers to the ability of different researchers to duplicate the same numerical results using the same techniques on the same set of non-experimental data.

⁹ Marcat and Sargent (1989), for instance, showed that the low inflation stationary equilibrium, π_{low} , would be stable and the high inflation stationary equilibrium, π_{high} , would be unstable if agents forecast inflation by running a least squares autoregression on past prices.

¹⁰ Similarly, Arifovic (1995) argues that the experimental findings are similar to those generated by a genetic algorithm—an adaptive, stochastic search process based on principles of population genetics, which also converges to a neighborhood of the low inflation steady state.

different government policies might have on subject behavior. They consider as an alternative to Equation 2 a monetary policy of the following form:

$$(3) \quad \Delta m = m_t - m_{t-1} = (\pi^* - 1)m_{t-1},$$

where π^* is the government's constant *gross inflation target*.¹¹ With Equation 3 in place of the constant government expenditure rule, (Equation 2), there is again the possibility of two stationary equilibrium values for inflation. The lower of these two stationary inflation rates is π^* , the government's targeted value, and the higher stationary inflation rate, π^{**} , is one in which fiat money ceases to have value and individuals merely consume their endowments. As in Equation 2, under rational expectations, the higher stationary inflation rate, π^{**} , is predicted as the limiting inflation rate for most initial conditions.

Marimon and Sunder (1994) report that when the policy rule, Equation 3 is used in place of Equation 2, and the government's target, π^* , is common knowledge, the inflation rate generated by the experimental subjects' decisions is again in a neighborhood of the low-inflation stationary equilibrium—in this case, π^* . They also considered whether preannounced changes in the two policy variables, d or π^* , have any effect on the decisions of experimental subjects. The rational expectations prediction is that a decrease (increase) in either d or π^* leads to a decrease (increase) in the low-inflation stationary equilibrium value for inflation (π_{low} or π^*). Marimon and Sunder report that preannounced *decreases* in these policy parameters do not have much of an anticipatory effect on subject behavior. Nevertheless, the inflation rate generated in the laboratory experiments did gradually drift downward to a neighborhood of the new lower stationary equilibrium, and the timing of this change was coincident with the actual change in policy.

Given that the two policy rules, Equation 2 and 3, seem to lead to the same experimental result, namely coordination near a stationary, low-inflation equilibrium, a natural question that arises is which of

Figure 2

Inflation Dynamics in the Model with a Growing Stock of Fiat Money


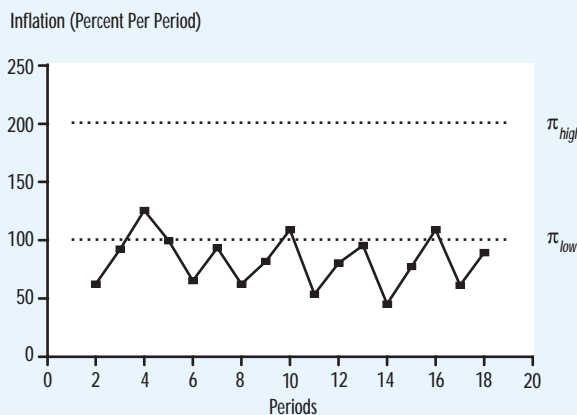


Figure 3

The Path of Inflation in Marimon and Sunder's (1993) Experimental Session 7C



these two rules is preferable from a policy perspective. For example, which rule leads to lower volatility in inflation? This is the kind of question that would be difficult to ask outside of a controlled laboratory setting, as natural experiments that could be used to address such questions are rather difficult to arrange. And even if such natural experiments could be arranged, it would be difficult to disentangle the public's reaction from issues surrounding the government's commitment to its new policy. In the laboratory, however, the policy rule is under the direct control of a non-elected official (the experimenter)—a fact that serves to

¹¹ Under this constant money growth rule, government revenue is now endogenously determined according to $g_t = ((\pi^* - 1) / \pi^*) \cdot m_{t-1} / p_{t-1}$.

minimize, though perhaps not eliminate credibility issues. Such control makes it a relatively easy task to consider the performance of different rules. The issue of whether the policy rule described by Equation 2 or by Equation 3 leads to lower inflation volatility is addressed in another paper by Marimon and Sunder (1995).

In this experiment, Marimon and Sunder considered a parameterization of the overlapping generations model where the inflation target, π^* , under the constant money growth rule, Equation 3, is nearly the same as the low-inflation steady state value, π_{low} , under the constant government expenditure rule, Equation 2. They then asked whether one of the two announced policy rules is better at stabilizing inflation. In this set of experiments, Marimon and Sunder (1995) simplified their experimental design considerably so that young agents only provided an inflation forecast (along with nonparticipating agents). This inflation forecast was used by the computer program to determine the optimal quantity of chips that each young agent sells for francs. Thus the game is now one in which optimization decisions are entirely absent and agents are simply concerned with forecasting the one-step-ahead inflation rate as closely as possible.

Though Marimon and Sunder (1995) claim, on the basis of their experimental evidence, that there is relatively lower volatility in inflation under the simpler constant money supply growth rule, Equation 3, than under the less transparent constant government expenditure rule, Equation 2, the evidence they present does not clearly demonstrate any significant difference. Under both policy rules, Marimon and Sunder report a few experimental sessions in which inflation rates persistently cycle around the low-inflation steady state, even though cyclic behavior for inflation does not constitute a rational expectations equilibrium.¹² They show that the cyclical behavior of inflation in these cases is tied to the adaptive forecast processes that agents appear to be using.¹³ Moreover, they do not find that inflation forecasts differ much under the two policy regimes. In particular,

there is no tendency for increasing coordination of inflation forecasts around the announced inflation target value, π^* .

This evidence might be taken to suggest that a policy of inflation targeting is not any better at coordinating inflationary expectations than constant money growth rules. Indeed, because inflation-targeting regimes are relatively new, there is little in the way of nonexperimental, empirical evidence to suggest that inflation-targeting policies are any better at coordinating inflationary expectations than earlier policy regimes. See, for example, Mishkin and Posen (1997). Hence experimental evidence of the sort presented by Marimon and Sunder might be useful in these sorts of policy discussions.

In extrapolating from experimental results to aggregate macroeconomic activity, however, much caution seems warranted. First, the amount of time that subjects have to learn about policies in the laboratory is not really comparable to the amount of time the general public would have available to react to these same policies. Experimental economists seek to lessen this problem by having subjects play many rounds of the same game, in an attempt to compress the amount of time it takes them to learn. There is no guarantee, however, that this compression of learning is a good approximation of the learning process that occurs outside the laboratory. Second, the policy rules themselves are highly simplified characterizations (some would say caricatures) of actual monetary regimes. For instance, under either policy rule, Equation 2 or 3, all government expenditures are financed through seigniorage, a rather unrealistic assumption. Still, Marimon and Sunder are only testing models that theorists themselves have used to characterize monetary policy regimes. Following a grand tradition in economic research, these models (and experimental results from tests of these models) are best viewed as abstract versions of reality.¹⁴

In the overlapping generations experiments reported thus far, explosive paths for inflation have not been observed. However, there is no disputing the fact that,

¹² Marimon and Sunder rule out the possibility that agents have coordinated on some kind of stationary sunspot equilibrium because they did not include a sunspot variable (a purely extrinsic variable) for individuals to condition on. In Marimon, Spear, and Sunder (1993), a two-state sunspot variable was introduced into an experimental overlapping generations economy. Realizations of this variable were not found to play a role in subjects' expectations. However, when the realizations of the sunspot variable were correlated with real shocks for awhile, the sunspot variable did appear to be a factor in subject's expectations even after the real shocks had ended.

¹³ In particular, they find evidence of learning equilibrium as described in Bullard (1994).

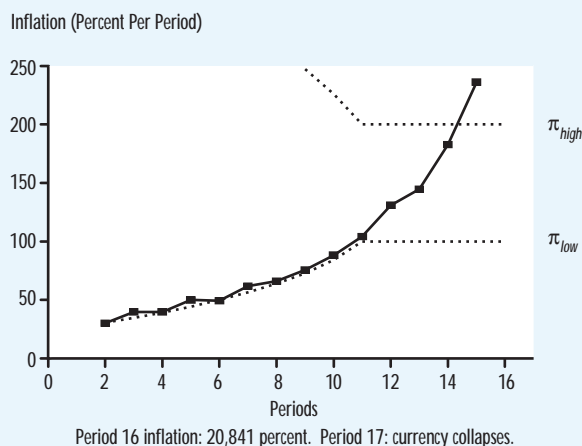
¹⁴ Other overlapping generations experiments, which test the effectiveness of different monetary policy regimes, include Bernasconi and Kirchkamp (1997), who test whether the labeling of policy regimes matters and Evans, Honkaphija, and Marimon (1998), who consider whether constitutional restraints on the amount of government seigniorage can be used to lower the stationary equilibrium value of inflation. Arifovic (1996) implements an experimental overlapping generations economy with two currencies and examines the behavior of the endogenously determined exchanged rate.

though infrequent, there have been historical episodes of hyperinflation and currency collapse. How might such hyperinflationary paths be generated in the laboratory? Ochs (1995) suggested that rather than adhering to just one or two constant money growth rates or expenditure levels, what may be required to generate hyperinflationary paths is a policy regime in which the growth rate of the money supply accelerates for some period. Here we have an example of historical evidence (on the likely source of hyperinflationary episodes) affecting experimental design. Following up on Ochs's suggestion, Hazlett and Kernen (1997) conducted overlapping generations experiments using the Marimon and Sunder (1995) experimental design, in which they steadily increased the value of government expenditures, d , over the first 10 periods of an experimental session and then held d constant. This has the effect of increasing the low stationary equilibrium value for inflation, but decreasing the high stationary equilibrium value for inflation, before both stationary values level off. Subjects were informed in advance of the number of chips (d) that the government would buy in every period. Though the level of d always remained below the maximum sustainable level, the steady increase in d did lead to a hyperinflationary path for prices and a currency collapse in two out of three sessions with inexperienced subjects. One of these two sessions is illustrated in Figure 4.

We see that in this session, actual inflation tracks the stationary low-inflation equilibrium rather well through Period 11. When the government ceases to increase its chip consumption after Period 11 (so that d is constant), inflation nevertheless continues to increase. After Period 16, the currency collapses, as young players' inflation forecasts were so high that it was optimal for them not to exchange any of their consumption good for francs. Interestingly, Hazlett and Kernen found that the currency collapse outcome was possible only with inexperienced subjects. When they brought experienced subjects back to play the same game, they found that the experienced subjects always coordinated their expecta-

Figure 4

The Path of Inflation in Hazlett and Kernen's (1997) Experimental Session 2B



tions around the low-inflation stationary value for inflation and tracked this value rather closely, even after government expenditures, d , had stopped increasing. This last finding suggests that hyperinflationary currency collapses may be sobering affairs, having a lasting impact on the way individuals evaluate government policies and update their expectations. One can also interpret this finding as support for the notion that agents eventually learn to form rational expectations in the sense that they avoid making systematic mistakes.

In all of the overlapping generations experiments where two stationary, rational expectations equilibria coexist, it is always reported that either individuals coordinate on the stationary equilibrium with the lower inflation rate or there is an explosive hyperinflation and currency collapse. Coordination on the high-inflation stationary rational expectations equilibrium never occurs. One might ask why the stationary, high-inflation equilibrium is ever focal under any expectations scheme, as it is a Pareto inferior outcome and is associated with counterfactual policy prescriptions: A reduction in government spending, d , or the money growth rate, π^* , leads to an

increase in the stationary value of π_{high} or π^{**} . Experimental evidence suggests an alternative possibility: The lack of attractiveness of the high inflation stationary equilibrium is caused by individuals' use of some kind of adaptive learning process as opposed to rational expectations. That is, under a variety of adaptive expectations schemes, the low inflation stationary equilibrium is locally stable and the high-inflation stationary equilibrium is locally unstable. Under this interpretation, the classical policy prescription within the overlapping generations framework—that a reduction in the money growth rate is associated with lower inflation—is one that presumes that agents learn over time in an adaptive manner.

Money as a Medium of Exchange

In McCabe's exchange game and in the overlapping generations experiments, the primary role of money is as a store of value. If, as has been shown, the money good serves as a store of value, then it also serves as a medium of exchange because the money good is the only means of storing wealth from one period to the next in all of these economies.¹⁵ However, as Ostroy and Starr (1990) have pointed out, though goods that serve as media of exchange must necessarily be stores of value, the reverse is not true—stores of value are not necessarily media of exchange. For example, old master paintings or baseball cards might be excellent stores of value but adopting them as commodity monies would require that these goods be conventionally used in trade. To understand the conditions under which a store of value comes to serve as a medium of exchange, what is needed is a model with multiple competing stores of value along with some motive for individuals to engage in trade. Kiyotaki and Wright (1989) have proposed such a model and have characterized the conditions under which certain stores of value may come to serve as media of exchange. Before we can discuss experimental tests of this model, it will be useful to have a brief sketch of the search-

Table 1

Each Type's Consumption and Production Goods

	Player Type		
	1	2	3
Desire to Consume Good:	1	2	3
But Produces Good:	2	3	1

theoretic model of money as a medium of exchange and the predictions that emerge from this theory.

In Kiyotaki and Wright's (1989) model, a population of agents is divided equally into three different types, labeled as Types 1, 2, or 3. In addition, there are three different goods, labeled as Goods 1, 2, or 3. Each player, Type i , desires to consume the good corresponding to his type, Good i , but produces some other good, $j \neq i$. Suppose that the distribution of production goods over types is as given in Table 1. Notice that, by design, in the absence of trade there is never a *double coincidence of wants* between any two players; that is, no agent produces the good desired by the player type who produces his consumption good. For trade to occur, some individuals must accept a good they do not want to consume with the expectation that they will be able to use this good to obtain a good they do want to consume. We will refer to a good used for such exchange purposes as a *medium of exchange*. Agents in this model are able to store a single unit of any good in any period, but storage of goods is costly. Kiyotaki and Wright suppose that the storage costs of Goods $i = 1, 2, 3$, denoted by c_i , are such that:

$$0 < c_1 < c_2 < c_3 < u,$$

where u is the utility that all three types get from consuming the good corresponding to their type. In every period, agents in this model are paired randomly with one another. The decision they face is whether to trade the good they currently have in storage for the good the agent with whom they are paired has in storage. Trades must be mutually

¹⁵ Similarly, in the experimental economy studied by Lian and Plott (1998), a cash-in-advance requirement for purchases ensures that money serves as a medium of exchange.

agreed on and involve simple one-for-one swaps of goods in inventory. If a player successfully trades for his consumption good, he immediately consumes that good and produces a new unit of his production good, which becomes the good he has in storage. In this case the player earns a positive net payoff equal to $u - c_j$, where c_j is the storage cost of his production good. If trade is not mutually agreed on or a player trades for a good that is not his consumption good, his net payoff for the period is $-c_j$, where c_j is the storage cost of the good he has in storage at the end of the period.

In the theoretical model, each player's trading decision involves weighing the cost of not trading and incurring the cost of the good currently held in storage against the expected net utility benefit from trading for the good held by the other player. If the other player's good is the desired consumption good, then, by design, it is always optimal to offer to trade for this good. The more interesting trading situations are those in which a player meets a good in trade that is not his consumption good and that is different from the good he currently has in storage. In such cases, the expected benefits from trade involve an analysis of the relative *marketability* of the good the player currently has in storage and the good that his trading partner has in storage. These expected benefits must be weighed against the difference in storage costs between the two goods.

Kiyotaki and Wright focus on steady-state, pure-strategy Nash equilibria. They show that for all valid parameterizations of this model, Type 2 and Type 3 players maximize their utility by adhering to fundamental, or cost-reducing, pure strategies in which they agree to trade for goods other than their consumption good only if these goods have lower storage costs than the good they currently have in storage. Thus Type 3 players will never offer to trade their production Good 1 for any good other than Good 3, because a trade of Good 1 for Good 2 would increase their storage costs. On the other hand, Type 2 players will always offer to trade their production Good 3 for Good 1, because this trade lowers Type 2's storage costs.

The optimal trading strategy of Type 1 players depends on parameter conditions. Type 1 players play the fundamental strategy of refusing to trade their production Good 2 for Good 3, which is more costly to store, provided that

$$(4) \quad c_3 - c_2 > (p_{31} - p_{21}) \frac{\beta u}{3}.$$

Here, p_{i1} denotes the fraction of agents of Type $i = 2, 3$ holding Good 1; $\beta \in (0, 1)$ is the constant discount factor. The right hand side of this inequality represents the discounted, expected difference in utility from having Good 3, rather than Good 2 in storage. The left hand side, $(c_3 - c_2)$, denotes the additional storage cost of Good 3 over Good 2. If this storage cost difference outweighs the expected gain in utility, then it is optimal for Type 1 players to play fundamentally and refuse to trade Good 2 for Good 3.

Provided the inequality in Equation 4 holds, the unique equilibrium prediction is that all three types play fundamental strategies. The pattern of exchange in this fundamental equilibrium is depicted in Figure 5. Type 2 players trade Good 3 to Type 3 players in exchange for Good 1. Type 2 players then trade Good 1 to Type 1 players in exchange for Good 2. Because Type 2 players do not value Good 1 for consumption purposes, but nevertheless agree to trade for it, Good 1 may be regarded as a *commodity medium of exchange* in this environment.

If the inequality in Equation 4 is reversed, there is a different unique pure-strategy Nash equilibrium in which Types 2 and 3 adhere to fundamental trading strategies but Type 1 players choose a speculative trading strategy, offering to trade their production Good 2 to Type 2 players in exchange for the more costly-to-store Good 3. Type 1 players then trade Good 3 with Type 3 players in exchange for Good 1, which Type 3 players produce. (See the illustration of this speculative equilibrium in Figure 6.) In this speculative equilibrium, there are two commodity media of exchange, Goods 1 and 3. Furthermore, because Good 3 is more costly to

Figure 5

The Pattern of Exchange in the Fundamental Equilibrium





Figure 6

The Pattern of Exchange in the Speculative Equilibrium



store than Good 1, Good 3 is dominated in rate of return as a medium of exchange by Good 1. The notion that a medium of exchange or money may be *dominated in rate of return* by other assets (in this case by another medium of exchange) is a real-world feature that is captured by Kiyotaki and Wright's model but is not so easily obtained in other models of money.

Kiyotaki and Wright (1989) provide only a characterization of the kinds of equilibria that are possible in their model. An important but unanswered question is whether these equilibria might ever be achieved by agents who do not begin a process of social interaction with equilibrium

beliefs as the theory presumes, but who must nevertheless adjust their strategies to their evolving historical experiences within the given structure of the model. Several experimental studies of the Kiyotaki-Wright environment have been conducted by Brown (1996) and Duffy and Ochs (1998a and 1998b) in an effort to address this question.

Brown (1996) considered a parameterization of the Kiyotaki-Wright model in which the unique Nash equilibrium prediction was for Type 1 players to speculate in the high cost Good 3, but for Type 2 and 3 players to play fundamental strategies, as illustrated by the pattern of exchange shown in Figure 6. Brown conducted two experiments in this environment in which one-third of the pool of 18 subjects were assigned the role of Type 1, one-third the role of Type 2, and one-third the role of Type 3. Subjects were informed of each type's objectives, told which good they currently had in storage, and then asked to submit trading strategies for all possible encounters with other agents. That is, would they offer to trade the good they currently had in storage if they met Type i with Good $j \neq i$ in storage? Subjects were then paired randomly and their trading strategies were implemented. They were informed of the results of these trading encounters—for instance, whether the good they had in storage had changed—and they were also informed of the dollar amount they had won or lost for the round. They then would have the opportunity to change their trading strategies before the next round. Subjects played approximately 50 rounds, with the end to a session determined randomly.

Brown reports that nearly all subjects offered to trade for their consumption good when given the opportunity. Furthermore, 99 percent of Type 2s offered to trade their production Good 3 for the less costly-to-store Good 1 when faced with the opportunity. Hence, Good 1 served as a medium of exchange in the experimental economy. Brown also found, however, that only 31 percent of Type 1 players offered to trade their production Good 2 for the more costly-to-store Good 3, even though the parameterization of the model was such


that Type 1 players' equilibrium best response was always to agree to such trades.

Duffy and Ochs (1998a) also tested the previously described version of the Kiyotaki-Wright model in a laboratory setting but sought to provide subjects with more information than was available in the Brown study. In particular, they provided subjects with population-wide information on the distribution of goods held in storage by each player type. They thought this information would help subjects assess the relative marketability of the goods they might meet in trade. In the theory, such information is common knowledge. Moreover, Duffy and Ochs allowed subjects to observe the type of player they were matched with and the good that this player type had in storage, before deciding whether to trade the good they currently had in storage. Finally, these authors made some effort to implement discounting of future payoffs, as the theory assumes, and to control for subjects' different attitudes toward risk. Subjects began a game with 100 points and gained or lost points based on the outcome of their trading decisions using known storage costs and other parameter values. Subjects' probability of winning a high cash prize was a linear function of their end-of-game point totals.

Duffy and Ochs considered two different parameterizations of the model—in particular, two different values for u , the utility value of consumption. One value was consistent with the fundamental equilibrium as depicted in Figure 5 and the other with the speculative equilibrium as depicted in Figure 6. For each parameterization, they report the results of five experimental sessions, each involving 24 to 30 subjects who participated in an average of 100 trading rounds. Like the Brown (1996) study, Duffy and Ochs found that nearly all subjects offered to trade for their consumption good when given the opportunity. Some of their other main findings are summarized in Figure 7, which reports average aggregate offer frequencies from five sessions in which the fundamental equilibrium was the unique Nash equilibrium and five sessions in which the speculative equilibrium was the unique Nash equilibrium.

Figure 7

Aggregate Results From Duffy and Ochs (1998a)



In the parameterization consistent with the speculative equilibrium, they found results that were similar to those of Brown (1996). An average of 93 percent of Type 2 players offered to trade Good 3 for Good 1, whereas an average of only 36 percent of Type 1 players chose to speculate by offering to trade Good 2 for Good 3. When the parameterization of the model was changed so that the unique Nash equilibrium was the fundamental equilibrium, Duffy and Ochs found that these aggregate offer frequencies were not significantly different. Nearly all Type 2 players, 97 percent, continued to offer to trade Good 3 for Good 1, as predicted in the fundamental equilibrium. Thirty percent of Type 1 players, however, speculated in Good 3, despite the fact that such speculation did not constitute a steady state best response. Duffy and Ochs attempted to uncover the motivations for players' actions (especially those of Type 1 players). Using a simple learning model, they showed that subjects were motivated mainly by the payoff success of past actions and not by the marketability concerns that the theory emphasizes.

Duffy and Ochs (1998b) considered a modification to the Kiyotaki-Wright environment in which an intrinsically worthless, but perfectly durable, fiat object, Good 0, was added to the model. This Good 0 could not be produced, consumed, or discarded by any player type. The fiat object was introduced into the economy by endowing a fraction, m , of each player type with Good 0 and endowing the remaining fraction, $1 - m$, with their production good. In one of the experimental treatments, the storage cost of Good 0 was set at zero, making it the least costly-to-store good; the storage costs of the other goods were unchanged. In this case there is an equilibrium in which Good 0 is always accepted when offered in trade. There is another equilibrium, however, in which no one believes that Good 0 has any value. This belief becomes self-fulfilling because no agent offers to trade for Good 0. In either case, Good 1 continues to circulate as a commodity medium of exchange in the fundamental equilibrium of the model.

Thus agents in this environment face an interesting coordination problem: Whether to avoid the use of an intrinsically worthless object as a medium of exchange in favor of a valued commodity object or to use both as media of exchange. Duffy and Ochs (1998b) report that in the experimental economy there is, in fact, a substantial volume of trade in the fiat object, Good 0, which is accepted anywhere from 80 percent to 95 percent of the time. Moreover, Good 1 remains a medium of exchange as well.

The findings that emerge from these experimental studies of money search models provide somewhat mixed support for the theory. On the one hand, goods may serve as media of exchange if they have relatively low storage costs compared with other goods. Moreover, like fiat money, these low-storage-cost goods need not be intrinsically valued. These findings are consistent with equilibria in which all players play according to fundamental, cost-reducing strategies. However, in parameterizations in which some players are predicted to play speculative trading strategies in equilibrium (Type 1 players in the version of the model described here) using relatively high-storage-cost goods as media of exchange, the experimental evidence appears at odds with these theoretical predictions. Further experiments in search-theoretic environments might help resolve these issues. On the other hand, these experimental findings might serve to motivate further theoretical explanations for why money is held despite being dominated in rate of return by other assets.

Money as a Unit of Account

Finally, we consider the role of money as a unit of account. There is no question that money serves as a unit of account, that is, as a common measure of the value of all goods and services. Indeed, most of us do not hesitate to think of the value of all goods and services in terms of the number of money units it takes to purchase them—that is, in terms of prices. However, when it comes to thinking about the value of goods and services over time, taking infla-

tion (or deflation) of prices into account, casual empiricism suggests that individuals frequently think in *nominal*, or current-dollar, terms, rather than in *real*, or constant-dollar, terms as is assumed in economic theories.

For example, the U.S. financial press is fond of reporting on changes in the benchmark Dow Jones Industrial Average (DJIA), which is the current-dollar, weighted average price of 30 Blue Chip stocks. If there is inflation of the price level then, over time, the value of the DJIA will reflect this inflation as well. What stock market investors should care about is changes in the value of the DJIA in constant dollars, that is, in dollars adjusted for inflation (as measured by changes in a price index). But that is not the figure that gets reported.

Irving Fisher (1928) defined money illusion as the “failure to perceive that the dollar, or any other unit of money, expands or shrinks in value.” Fisher argued that money illusion was a serious problem because economic contracts are nearly always specified in money units, and the parties to such contracts must recognize the possibility of inflation or deflation of prices over the length of the contract to properly judge whether the contract is likely to result in a real gain or loss. Because Fisher thought that money illusion was pervasive, he argued that it was the task of the monetary authority to stabilize the level of prices and therefore the purchasing power of money. He likened this task to government efforts to standardize weights and measures to prevent *cheating*. “If,” Fisher argued, “we are at such pains to standardize or stabilize the yard, the pound and every other goods-unit, having—as we do—official settlers of weights and measures to prevent the cheating of the public because of changes in those goods-units, how much more important is it to stabilize the unit of money, applying as it does to every purchase and sale.”

Fisher’s argument for price stability is predicated on the notion that the public suffers from money illusion, but what is the hard evidence for this claim? Fisher himself related conversations he had with Germans he met while touring Germany,

as well as with various Americans, (including his own dentist!) to buttress his contention that money illusion was pervasive. Since Fisher’s (1928) book, however, economists have largely ignored the empirical question of whether individuals suffer from money illusion, preferring to believe that there can be no such illusion. As Howitt (1987) notes, money illusion contradicts economists’ assumptions of rational, utility, or profit-maximizing agents, and for this reason economists tend to dismiss the possibility of money illusion out of hand.

Recently, the question of the extent to which individuals may suffer from money illusion has been reconsidered by Shafir, Diamond, and Tversky (1997). Going beyond the casual empiricism of Fisher and others, these authors have sought to document, through the use of survey questionnaires, the extent to which individuals think in nominal rather than real terms.¹⁶ Though the use of survey questions is prone to a number of biases and differs in many respects from controlled laboratory environments where individuals are paid on the basis of the decisions that they make, the idea of attempting to quantify the extent of money illusion by asking the public how they think about economic decisions is similar to the experimental methodology.

Shafir et al. used a variety of different survey questions that were designed to detect the presence or absence of money illusion in economic decision-making. The questions were presented to large groups of people (typically 100 or more) at two New Jersey shopping malls, the Newark International Airport, and to Princeton undergraduates. In most cases, the answers given by the various groups were similar. The authors conclude that the answers to these survey questions indicate that “money illusion is a widespread phenomenon in the United States.”

Consider for example, the following question, which is typical of the survey questions in the Shafir et al. study:

“Two competing bookstores have in stock an identical leather-bound edition of

¹⁶ Similarly, Bewley (1998) uses surveys of firm behavior to examine the extent to which labor markets suffer from a related illusion, namely resistance to nominal wage cuts. Shiller (1997a) uses survey evidence to explain why there is resistance to indexation of long-term contracts. In another paper, Shiller (1997b) uses survey evidence to demonstrate that economists’ mental models of inflation are substantially different from those of non-economists.

Oscar Wilde's collected writings. Store A bought its copies for \$20 each. Tom, who works for Store A, has just sold 100 copies of the book to a local high school for \$44 a copy. Store B bought its copies a year after Store A. Because of a 10% yearly inflation, Store B paid \$22 per copy. Joe, who works for Store B, has just sold 100 copies of the book to another school for \$45 a copy. Who do you think made a better deal selling the books, Tom or Joe?"

Shafir et al. report that 87 percent of a group of 130 subjects incorrectly chose Tom. The authors hypothesize that these subjects focused on the fact that Tom's nominal profit of \$24 was higher than Joe's nominal profit of \$23, even though in real terms Joe earned more than Tom.

Some might argue that findings of this sort are simply a result of the manner in which the problem is framed, or presented, to subjects. Indeed, Shafir et al. address the possibility of framing effects at some length. The framing problem—that results may differ under different frames—is one that plagues all research, not just survey or experimental methods. Indeed, it is difficult to imagine developing any model of behavior that does not rest on some particular framing of a problem.

To illustrate the potential importance of framing effects, suppose the wording of the preceding question was changed as follows:

"Tom and Joe both paid the same real price for 100 copies of an identical leather-bound edition of Oscar Wilde's collected writings. Tom has just sold his 100 copies to a local high school for \$44 while Joe has just sold his 100 copies to another school for \$45. Who do you think made the better deal selling the books, Tom or Joe?"

With this new wording, I would hazard that most would find the correct answer (Joe) rather obvious. The new wording of the problem, however, is such that there is no possibility for money illusion to occur—that is, the problem has been assumed away, as in most economic

theorizing. As Shafir et al. point out, there are ways of providing more neutral frames, which eschew terms such as *real*. Using these more neutral frames, the authors report that the incidence of money illusion is only somewhat diminished.

A stronger critique of Shafir et al.'s survey results is that the decisions subjects were asked to make were, in all cases, hypothetical, and with one exception, subjects were not paid according to the answers they provided.¹⁷ In fact, most subjects were not paid anything to answer the survey questions, which might make us wonder whether they were sufficiently motivated to think hard about their responses. It thus remains to be seen whether these same kinds of results will continue to be obtained when an individual's decisions have direct and significant payoff consequences as they do in actual markets.¹⁸ Indeed, one can think of the survey findings as a kind of pilot test for controlled laboratory experiments; if the survey results did not reveal evidence of money illusion, then it seems unlikely that paying subjects according to their decisions would lead to significantly different results. Casual empiricism of the type alluded to earlier suggests that it should be possible to find evidence of money illusion with paid human subjects in controlled laboratory settings.¹⁹

CONCLUSION

What, if anything, can we learn from these experimental studies? Can experimental evidence be used to confirm or refute a particular monetary model or equilibrium prediction? Some might argue that the theoretical environments examined, especially the overlapping generations and search theoretic models of money, are too complex for subjects to comprehend, and for this reason, these theories cannot be tested in the laboratory. This argument, however, begs the question of just what type of behavior these theories seek to describe. If a theory cannot be tested, because of a lack of non-experimental data and because laboratory tests of it are judged to be impossible, then what one has, essentially, is a vacuous theory.

¹⁷ The exception is an experiment conducted by Thaler, Tversky, Kahneman, and Schwartz (1997) which is discussed in Shafir et al. (1997).

¹⁸ Tversky and other psychologists have argued that in many (but not all) cases, subjects' responses to hypothetical survey questions do not depend on whether the subjects are paid according to their responses. Experimental economists, for example, Smith and Walker (1993), have disputed this view, with experiments that show that paying subjects according to the decisions they make does make a big difference in behavior and nearly always reduces the variance of observed actions.

¹⁹ See, for example, the work of Fehr and Tyran (1998).

A more reasonable critique is that the experimental environment is not complex enough or that it leaves out many real world features—for example, communication or observation of the decisions of others, which, if allowed, would serve to further promote the adoption of money as a store of value, as a medium of exchange, or as a unit of account in accordance with the theoretical predictions. But this same critique can be made of the theory itself, which does not allow for, or more typically has nothing to say about, such real-world features and is itself only an approximation to some observed real-world phenomena.

The approach of experimental researchers is to attempt to approximate the theoretical environment as closely as possible and then argue that the experimental results from the approximated environment can be used to test the theory. In many cases, as we have seen, these laboratory approximations deliver results that are consistent with some of the theoretical predictions. Furthermore, in the overlapping generations and search-theoretic environments where multiple equilibria are possible, evidence from laboratory experiments can be used to resolve the question of which equilibrium is likely to be more focal, and therefore more relevant, for policy purposes or theoretical extensions.

In cases where the experimental results are at odds with the theory, there is always the possibility that further refinements of the experimental environment—for instance, changes in the manner in which problems are framed—may serve to deliver experimental results that ultimately confirm the theoretical predictions. There should also, however, be feedback from experimental results to theoretical modeling, especially when experimental results are continually found at odds with a particular theory.

The experimental studies discussed in this paper provide some support for several different theories of money, but large gaps still remain in our understanding of why money is held and how decisions involving money are made. Further theoretical work, in combination with laboratory experiments can serve only to deepen our understanding of this important topic.

REFERENCES

- Arifovic, Jasmina. "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economics," *Journal of Political Economy* (June 1996), pp. 510-41.
- _____. "Genetic Algorithms and Inflationary Economies," *Journal of Monetary Economics* (December 1995), pp. 219-43.
- Bernasconi, Michele, and Oliver Kirchkamp. "The Name of Inflation: An Experimental Study of Saving, Inflation and Monetary Policies in an Overlapping Generations Economy," *Universitat Mannheim Working Paper* (September 1997).
- Bewley, Truman F. "Listening to Business: A Study of Wage Rigidity," *Yale University Unpublished Manuscript* (January 1998).
- Brown, Paul M. "Experimental Evidence on Money as a Medium of Exchange," *Journal of Economic Dynamics and Control* (April 1996), pp. 583-600.
- Bullard, James. "Learning Equilibria," *Journal of Economic Theory* (December 1994), pp. 468-85.
- Cagan, Phillip. "The Monetary Dynamics of Hyperinflation," *Studies in the Quantity Theory of Money*, Milton Friedman, ed., University of Chicago Press, 1956, pp. 25-117.
- Dewald, William G., Jerry G. Thursby, and Richard G. Anderson. "Replication in Empirical Economics: The Journal of Money, Credit and Banking Project," *American Economic Review* (September 1986), pp. 587-603.
- Duffy, John and Jack Ochs. "Emergence of Money as a Medium of Exchange: Experimental Evidence," *University of Pittsburgh Working Paper* (January 1998a) and *American Economic Review*, forthcoming.
- _____. and _____. "Fiat Money as a Medium of Exchange: Experimental Evidence," *University of Pittsburgh Working Paper* (March 1998b).
- Evans, George W., Seppo Honkapohja, and Ramon Marimon. "Fiscal Constraints and Monetary Stability," *University of Oregon Working Paper* 217, 1998.
- Fehr, Ernst and Jean-Robert Tyran, "Does Money Illusion Matter? An Experimental Approach," *Discussion Paper* No. 9802, (February 1998), University of St. Gallen
- Fisher, Irving. "The Money Illusion", Adelphi, 1928.
- Hazlett, Denise, and Amanda Kernen. "Hyperinflation and Seigniorage in an Experimental Overlapping Generations Economy," *Whitman College Working Paper* (September 1997).
- Howitt, Peter. "Money Illusion," *The New Palgrave: A Dictionary of Economics*, vol. 3, John Eatwell et al., ed., The Macmillan Press, 1987, pp. 518-19.
- Kagel, John H., and Alvin E. Roth eds. "The Handbook of Experimental Economics," *Princeton University Press*, 1995.
- Kiyotaki, Nobuhiro, and Randall Wright. "On Money as a Medium of Exchange," *Journal of the Political Economy* (August 1989), pp. 927-54.
- Lian, Peng, and Charles R. Plott, "General Equilibrium, Markets, Macroeconomics and Money in a Laboratory Experimental Environment," *Economic Theory*, 12 (July 1998), pp. 21-75

- Lim, Suk S., Edward C. Prescott, and Shyam Sunder. "Stationary Solution to the Overlapping Generations Model of Fiat Money: Experimental Evidence," *Empirical Economics* 19(2), 1994, pp. 255-77.
- Lucas, Robert E. Jr. "Adaptive Behavior and Economic Theory," *Journal of Business* (October 1986), pp. S401-S426. Reprinted in *Rational Choice: the Contrast Between Economics and Psychology*, Robin Hogarth and Melvin Reder eds., Chicago: University of Chicago Press, 1986, pp. 217-42.
- Marcet, Albert, and Thomas J. Sargent. "Least-Squares Learning and the Dynamics of Hyperinflation," *Economic Complexity: Chaos, Sunspots, Bubbles, and Nonlinearity*, William A. Barnett et al., eds., Cambridge: Cambridge University Press, 1989, pp. 119-37.
- Marimon, Ramon. "Learning From Learning in Economics," *Advances in Economics and Econometrics: Theory and Applications*, vol. 1, David M. Kreps and Kenneth F. Wallis, eds., Cambridge University Press, 1997.
- Marimon, Ramon, Stephen Spear, and Shyam Sunder. "Expectationally Driven Market Volatility: An Experimental Study," *Journal of Economic Theory* (October 1993), pp. 74-103.
- Marimon, Ramon, and Shyam Sunder. "Expectations and Learning Under Alternative Monetary Regimes," *Economic Theory* (January 1994), pp. 131-62.
- _____ and _____. "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence," *Econometrica* (September 1993), pp. 1073-107.
- _____ and _____. "Does a Constant Money Growth Rule Help Stabilize Inflation? Experimental Evidence," *Carnegie Rochester Conference Series on Public Policy*, 1995, pp. 111-56.
- McCabe, Kevin A. "Fiat Money as a Store of Value in an Experimental Market," *Journal of Economic Behavior and Organization* (October 1989), pp. 215-31.
- McKelvey, Richard D., and Thomas R. Palfrey. "An Experimental Study of the Centipede Game," *Econometrica* (July 1992), pp. 803-36.
- Mishkin, Fredric S. and Adam S. Posen. "Inflation Targeting: Lessons From Four Countries," Federal Reserve Bank of New York *Economic Policy Review* (August 1997), pp. 9-110.
- Ochs, Jack. "Coordination Problems," *The Handbook of Experimental Economics*, John H. Kagel and Alvin E. Roth, eds., Princeton University Press, 1995, pp. 195-251.
- Ostroy, Joseph M., and Ross M. Starr. "The Transactions Role of Money," *Handbook of Monetary Economics*, vol. 1, Benjamin M. Friedman and Frank H. Hahn, eds., Elsevier Science Publishers, 1990, pp. 3-62.
- Sargent, Thomas J. *Bounded Rationality in Macroeconomics*, Oxford University Press, 1993.
- _____. "The Ends of Four Big Inflation," *Rational Expectations and Inflation*, Thomas J. Sargent, ed., New York: Harper and Row, 1986, pp. 40-109.
- Shafir, Eldar, Peter Diamond, and Amos Tversky. "Money Illusion," *Quarterly Journal of Economics* (May 1997), pp. 341-74.
- Shiller, Robert J. "Public Resistance to Indexation: A Puzzle," *Brookings Papers on Economic Activity* 1, (1997a), pp. 159-228.
- _____. "Why Do People Dislike Inflation?" *Reducing Inflation: Motivation and Strategy*, Christina Romer and David Romer, eds., University of Chicago Press (1997b), pp. 3-65.
- Smith, Vernon L., Gerry L. Suchanek and Arlington W. Williams. "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica* (September 1988), pp. 1119-51.
- Smith, Vernon L. and James M. Walker. "Monetary Rewards and Decision Cost in Experimental Economics," *Economic Inquiry* (April 1993), pp. 245-61.
- Thaler, Richard H., Amos Tversky, Daniel Kahneman, and Alan Schwartz. "The Effect of Loss Aversion on Risk Taking: An Experimental Test," *Quarterly Journal of Economics* (May 1997), pp. 647-61.
- Wallace, Neil. "A Dictum for Monetary Theory," Federal Reserve Bank of Minneapolis *Quarterly Review* (Winter 1998), pp. 20-26.