

Exploring Elements of Exchange Rate Theory in a Controlled Environment

Eric O’N. Fisher*
The Ohio State University
Federal Reserve Bank of Cleveland
+1.614.292.2009
fisher.244@osu.edu

First Draft: 23 July 2002
This Draft: 1 December 2003

Abstract

This paper analyzes the simple implications of exchange rate theories in the laboratory. It shows that purchasing power parity, covered interest parity, and uncovered interest parity fare well in simple environments. Not-traded goods and non-stationary domestic prices do cause deviations from these elements of exchange rate theories. But the experimental evidence indicates that non-stationary domestic prices have a much stronger effect than not-traded goods in causing deviations from purchasing power parity and the two interest parity conditions.

*The author thanks the National Science Foundation for grants SES-9870874 and SES-0111315 that made this work possible. He also is grateful for a Jean Monnet Fellowship from the European University Institute that allowed him to complete this project and the hospitality of Osaka University’s Institute for Social and Economic Research, where some of these ideas had their gestation. He thanks Charles B. Grant, David Hine, Philippe Andrade, and seminar participants at the European University Institute, The Johns Hopkins University, and the fall 2003 Economic Science Association Meetings for comments on earlier drafts. Namita Bafana and David Hine provided very able research assistance during this experiment.

1. Introduction

Perhaps the most celebrated article in the literature on exchange rate econometrics is Meese and Rogoff (1983). Those authors set the standard for an entire strand of literature because they had found that the impressive empirical findings of an earlier article by Frankel (1979), published in a leading journal of our discipline, broke down if one added two more months of data.¹ Meese and Rogoff's work has become seminal precisely because it was *not just* about empirical exchange rate models in the first decade of the modern era of floating rates. Those authors set the bar that a generation of scholars has tried to surpass, and twenty years later there is little evidence that exchange rate models perform better than a random walk when put to the task of predicting out of sample.² Why do these models perform so poorly?

This paper uses a decidedly unorthodox technique to explore this question. I employ the laboratory to look at the three most basic elements of exchange rate theory: purchasing power parity, covered interest parity, and uncovered interest parity. My research establishes two main points. First, each of these three elements of exchange rate theory fares well in the laboratory. Second, not-traded goods and non-stationary domestic prices do cause deviations from simple theoretical predictions, but a non-stationary environment has a more significant effect than does not-traded goods in explaining why simple predictions of exchange rate theories break down.

The entire literature on experimental economics and exchange rates consist of *four* papers: Arifovic (1996), Noussair, Plott, and Riezman (1997), Fisher and Kelly

¹ Please forgive me if I betray a trust by reporting a private conversation with one of the authors while he and I were both employed at the central bank of the United States.

² I am not unaware of the continuing debate on whether the random walk is the best out-of-sample predictor for exchange rates. Using an innovative non-parametric approach, Nelson Mark (1995) breathed some life into monetary models of the exchange rate by showing that they may be able to forecast a small but significant part of the variation in exchange rates. Faust, Rogers, and Wright (2001) seem to indicate that Mark's results may be specific to the particular vintage of the data that he used. I found the summary of the current state of the art in Neeley and Sarno (2002) to be well written and illuminating indeed.

(2000), and Fisher (2001).³ Arifovic (1996) studies exchange rates in an overlapping generations model and explores the celebrated result that the exchange nominal rate is not determined in this class of models. Perhaps because it is notoriously difficult to implement such a model in the laboratory, she found mixed support at best for the theory. Noussair, Plott, and Riezman (1997) analyze a two-country model with a real side and two cash-in-advance constraints; again they found mixed evidence for simple elements of exchange rate theory.

Fisher and Kelly (2000) study essentially identical assets in a non-stationary environment and showed that cross-asset arbitrage held, even though every asset had a significant bubble. We interpreted this evidence as support for simple exchange rate theories; even though subjects do not perform backward induction, they do understand that asset prices are related, and thus the bubbles for the different assets are almost perfectly correlated. A shortcoming of this research may be that there is no explicit role for assets *qua* foreign exchange. Fisher (2001) rectified this deficiency by designing an experiment based upon two cash-in-advance constraints in sessions that were the precursor to this current manuscript. The main difference between Fisher (2001) and Noussair, Plott, and Riezman (1997) is that the former has a much simpler design because the experimenter makes the supply side of both the goods and foreign exchange markets.

This paper is the first attempt to explore systematically *why* exchange rate models might fail. Instead of looking at all the details of a particular model -- such as a two-country cash-in-advance model with production according to comparative advantage -- it uses a balanced design and examines a very simple model of exchange rates with only very minor variants across treatments. The small differences are

³ Let me put this number in perspective. A search of EconLit in July 2002 found 16992 hits for the

intended to highlight elements of exchange rate theory. Thus all the sessions have two cash-in-advance constraints and simple goods and foreign exchange markets, but some sessions (designed to explore covered interest parity) have interest-bearing assets and others (designed to explore uncovered interest parity) have both interest-bearing assets and random foreign inflation. The building blocks of any model of exchange rate theory are purchasing power parity and the two different interest parity conditions, and this experiment gets right to the heart of these elements. Also, I describe data from fifteen sessions, many more than is the norm in the nascent literature on experimental macroeconomics.⁴

How does my work contribute to the larger literature on empirical models of exchange rates? I am not dogmatic enough to assert that the simple treatments I describe below capture the complexity of the international monetary system. But I do hope that the reader will agree that experimental economics might well complement traditional research using applied econometrics. Thus the results that I describe in this paper ought to serve further to reinforce the notion that both not-traded goods and non-stationary prices contribute to the empirical failure of exchange rate models. But this experiment shows clearly that non-stationary prices have a stronger effect. I believe this is the main contribution of this paper, and it shows how experimental economics and traditional econometrics can reinforce one another.⁵ I also show that a non-stationary environment increases the average dispersion of the subjects' bids for foreign exchange, and I propose tentatively that this phenomenon occurs because of

keywords "exchange rates," 786 hits for "purchasing power parity," and 2035 hits for "money demand."

⁴ For example, Arifovic (1996) ran two sessions, and Noussair, Plott, and Riezman (1997) ran four.

⁵ Using an intuitive statistical decomposition, Engel (1999) showed that not-traded goods prices accounted for almost none of the variability of the U.S. bilateral real exchange rate against several major countries over a long period. He concluded that persistent differences in the relative prices of traded goods were the main reason for the volatility of real exchange rates during the modern period of floating rates.

the cognitive difficulty of making decisions in an environment where domestic prices are trending. Perhaps this is a second contribution of my work, and it is an avenue that bears further investigation.

The rest of this paper is structured as follows. The second section states the six hypotheses that are the foci of the research. The third section gives an extended discussion of the design; it is important that the reader try to put himself or herself into the mind of the subject to see what it would be like to participate in these sessions. The fourth section describes the procedures, and the fifth section derives the equilibrium in the foreign exchange market. It shows that there are several equilibria and emphasizes that the no-surplus equilibrium is the basis for the theoretical predictions for the exchange rate. The sixth section gives the experimental results. The heart of the empirical findings is in the diagrams that show the exchange rates for each of the sessions. As is customary in the literature on experimental economics, I use non-parametric statistics to substantiate (or falsify) the hypotheses that are at the center of this work. Finally, the seventh section is a brief conclusion and an exhortation for future research.

2. Hypotheses

Let e_t be the domestic currency price of a unit of foreign exchange at time t . Also, let p_t be the domestic currency price of the home good and p_t^* be foreign currency price of the import, both measured at time t . Of course, p_t and p_t^* are treatment variables, and much of the statistical analysis focuses on e_t from the different treatments. The simplest hypothesis is relative version of purchasing power parity, the notion that the real exchange rate $q_t \equiv e_t p_t^* / p_t$ is independent of time. We postulate:

HYPOTHESIS 1: *If all goods are traded, then the relative version of purchasing power parity holds.*

The domestic good and its foreign counterpart have induced valuations according to their redemption values. Let α_c be the (dollar) redemption value of the domestic good and α_b be that of the import.⁶ Then, if all goods are traded, the real exchange rate $q_t \equiv e_t p_t^* / p_t = \alpha_b / \alpha_c$ ought to be constant and equal to the marginal rate of substitution between domestic and foreign goods. In brief, we state:

HYPOTHESIS 2: *If all goods are traded, then the absolute version of purchasing power parity holds.*

Another fundamental relationship has to do with interest-bearing assets. Consider a world without commodity price uncertainty. Then, if the interest rate on foreign currency deposits is higher than that on domestic deposits, the nominal rate exchange rate ought to depreciate exactly to offset this differential. Forward contracts on foreign exchange are actually redundant assets, since the forward premium automatically reflects the interest differential when there is no inflation risk. Hence, the rate of depreciation of foreign exchange ought to equal the interest differential exactly, even in treatments where there is no explicit asset mimicking a forward contract. We examine:

HYPOTHESIS 3: *If all goods are traded and there is no uncertainty, then covered interest parity holds.*

The analog of covered interest parity is that the expected depreciation of the spot rate reflects differences in expected inflation between countries, even when an open position in foreign exchange is subject to inflation risk. In this case, forward

⁶ The reader who is not *au courant* in experimental economics might wonder now these values arise. Again, they are treatment variables used to induce the subjects' common utility function, which maps

contracts are redundant only if all agents are risk neutral, share the same information sets, and have common prior beliefs about all payoff-relevant events. Still, one can test whether expected depreciation of domestic currency reflects the real interest differential. It is important to emphasize here that the experimenter has a decided advantage over the econometrician because the experimenter actually designs the inflation process and explains it to all the subjects, creating common beliefs about this stochastic process. Thus the experimenter actually knows the real interest rate, and both nominal interest rates and expected inflation are in everyone's (common) information set. The econometrician almost always has to treat unobserved expectations about inflation -- an ineluctable part of real interest rates -- as a nuisance parameter.⁷

HYPOTHESIS 4: If all goods are traded, then uncovered interest parity holds.

These four hypotheses are the central elements of most theories of exchange rate determination. They formed the crux of my earlier preliminary experimental analysis (2001) of exchange rates in models with explicit cash-in-advance constraints, and that work showed strong initial confirmation of each of these hypotheses. But these aspects of exchange rate theory are decidedly *not* corroborated in analyses based on statistical inference from historical field data. So the salient unanswered question is: What actually causes departures from exchange rate theories?

One important candidate is that not-traded goods are a fundamental source of measurement error when the econometrician uses wholesale or consumer price indices to test models of exchange rate determination. The crux of the argument is that even if the law of one price is empirically plausible for a wide array of traded goods, it is

from home and foreign goods purchased into dollars earned. This aspect of the experiment will become clearer below.

⁷ An insightful discussion of the pitfalls inherent in trying to measure inflation expectations from field data is found in Flood and Garber (1981).

not appropriate to generalize from this law to the properties of aggregate price indices because the latter contain a large fraction of not-traded goods. Their prices need not move in lockstep with changes in the nominal exchange rate. Of course, a controlled environment gives the researcher real empirical traction. I am afforded the great luxury of designing treatments that obviate the difficulties inherent in measuring price indices with error.

HYPOTHESIS 5: The empirical implications of exchange rate theory break down because of not-traded goods.

Another very plausible candidate that may explain why exchange rate theories fare so poorly under econometric scrutiny is that the macroeconomic environment is not stationary. Nominal exchange rates are asset prices, and there are good theoretical reasons to model them as Martingales. Also, if price indices have unit roots, then the typical regressions used to test these theories no longer satisfy the assumptions of Gauss and Markov. These two observations have given rise to a large cottage industry in the last fifteen years applying co-integration tests to exchange rates and prices indices.⁸ In the treatments designed to explore the effect of a non-stationary environment, I actually let the domestic currency price of home goods have a deterministic time trend. This is perhaps the simplest way to introduce a non-stationary environment, even though there is no “stochastic trend” or unit root in domestic prices. If this deterministic trend causes deviations from the theoretical predictions of elements of exchange rate theories, then it is perhaps quite likely that there would be even larger effects if the trends were actually stochastic and there were a unit root in domestic prices.

⁸ I confess to being one of the guilty parties in Fisher and Park (1991).

HYPOTHESIS 6: The empirical implications of exchange rate theory break down because the economic environment is not stationary.

These six hypotheses are the bones upon which I will hang the meat of the empirical analysis. I hope that the reader is sympathetic enough to agree that they are important elements of any model of exchange rate determination, and I am sure that the reader is wise enough to see that these hypotheses are not exhaustive. Acknowledging quite frankly that this work is only the first step in what I hope will become a much larger literature, I ask for your kind indulgence and proceed now to describe in detail the experiment's design.

3. The Experimental Design

The most important element of the design is that domestic currency and foreign exchange are both intrinsically worthless; a money has value only because it can be used to purchase a good, which then can be redeemed for cash at the end of the session. This element of the design imposes a strong cash-in-advance constraint. Some might argue that this aspect of the designs is unrealistic because only a very small fraction of trading in foreign exchange markets in the field occurs in order to finance imports. On the other hand, imposing two simple cash-in-advance constraints is the norm in a large body of the theoretical literature in international monetary economics.

The experiment had a 5×3 design. Table 1 summarizes the fifteen sessions that were run. Its columns are the three different treatments. In every session in the first column, all goods were traded and the macroeconomic environment was stationary. The sessions in the second column allow for a non-linearity of payoffs that captures the macroeconomic effects of not-traded goods. This aspect of the design will be discussed in detail in Section 3.3 below. The third column allows for

non-stationary prices for domestic goods; in each of these sessions, the domestic currency price of the home good and the concomitant supply of liquidity rise in a known and pre-determined fashion.

Each row represents an element of exchange rate theory. The sessions in the first and second rows analyze a simple version of purchasing power parity. Those in the third row have interest-bearing foreign assets but no import price variability; since there is no risk in holding an uncovered position in foreign currency, these sessions examine covered interest parity. The fourth and fifth rows have sessions in which foreign exchange pays interest and there is import price uncertainty. Since holding an open position in foreign exchange is now risky, these sessions explore uncovered interest parity.

Table 1: Experimental Design				
		<i>Macroeconomic Environment</i>		
		Control	Not-Traded Goods	Non-Stationary Prices
<i>Element of Exchange Rate Theory</i>	Purchasing Power Parity	5 October 00	8 February 02	1 February 02
		10 October 00	15 February 02	1 March 02
	Covered Interest Parity	13 October 00	8 March 02	22 February 02
	Uncovered Interest Parity	20 October 00 24 January 02	30 May 02 7 June 02	10:00 15 March 02 14:00 15 March 02

Since the experimental design constitutes a two-way layout with five blocks and three treatments, it is easy to examine treatment effects.⁹ A block is a group of three sessions, one in each of the treatments. In particular, if one runs *six* sessions on purchasing power parity with two in each macroeconomic environment, then one has *two* blocks. Thus the first two blocks are the observations on purchasing power parity, and they form the first two rows of Table 1. The third block consists of the observations on covered interest parity, and they form the third row of Table 1. The

last two blocks are the observations on uncovered interest parity, and they constitute the fourth and fifth rows of Table 1. Of course, all these observations are independent because each session is conducted with different subjects drawn from a fairly homogeneous population. So a general test for a treatment effect has the interpretation that either not-traded goods or non-stationary domestic prices give rise to different outcomes for at least one of the three elements of exchange rate theory, purchasing power parity, covered interest parity, or uncovered interest parity.

3.1 Details of the Foreign Exchange Markets

In every treatment, the experimenter made the supply side of the foreign exchange market. Each subject was endowed with sufficient domestic currency in each period to satisfy his or her needs for liquidity. Domestic currency could buy domestic goods, but it was worthless in itself. It could also buy foreign exchange. Again, foreign currency was worthless but it could be used to buy imports. (Both domestic goods and imports had known redemption values; I will discuss the goods market in greater detail in the next subsection.) The domestic currency was called German marks, and the foreign currency was called French francs. The domestic good was called German chocolate cake, and the foreign good was called French bread. It may be somewhat less boring if I use these names occasionally in the general descriptions in the rest of this section.

The supply of foreign exchange was completely inelastic. In particular if there were k subjects, then $k - 2$ units of foreign exchange were for sale. The markets were conducted as third-price auctions. Thus each agent submitted a sealed bid for foreign exchange in each period, and the bids were put in rank order. The two lowest bids did not get to purchase foreign exchange in that period, and the remaining $k - 2$

⁹ See Hollander and Wolfe (1973), chapter 7.

bidders each got to purchase one French franc at the highest excluded bid.¹⁰ In each period in each session, there was an upper bound on the admissible bids that the subjects could submit for foreign exchange; that upper bound was almost always much larger than the theoretical value of foreign exchange. No one could bid more marks than he or she had on hand, and the number of marks was refreshed in each period or sub-period.

Why did I conduct the foreign exchange markets as sealed bid auctions? I wanted to use a very simple mechanism whose properties were well known to make the demand side of this market. Since the essence of this research has to do with the price of foreign exchange, I wanted to make sure that this market worked well. By offering a fixed supply of foreign exchange and by setting up the demand side as a third-price auction, I created an environment in which the most plausible equilibrium exchange rate corresponds exactly to the theory from international finance.

Foreign exchange markets in the field are over-the-counter markets where the major commercial banks make both sides of a transaction. Also, a forward transaction very often offsets activity in the spot market, and many spot transactions are really thus part of a foreign exchange swap. Further, a wide array of derivative securities, including foreign exchange futures and options on foreign exchange, influence the spot markets for currencies. The foreign exchange markets in this experiment have none of these subtleties, but they have the major advantage of being analytically simple. It was very easy for the subjects to learn quickly how these markets worked, and they could see readily that the demand for foreign exchange was based solely upon the value of the foreign goods that it could purchase.

¹⁰ Thus I used the price determination rule that gives the entire surplus to the buyers.

3.2 Interest-Bearing Assets

In the treatments having to do with covered interest parity and uncovered interest parity, foreign exchange deposits paid a known nominal interest. In all the sessions in the third, fourth, and fifth blocks of Table 1, each period was divided into two sub-periods. Deposits in domestic currency always paid a zero nominal interest rate, but foreign exchange deposits paid a known nominal interest of 0%, 100% or 200% between sub-period A and sub-period B. This meant that a subject could buy foreign exchange in sub-period A, knowing that it would pay interest if held through sub-period B. In the three sessions in the third block, there was no variability in the foreign price of imports. In the treatments in the fourth and fifth blocks, imports cost one or two units of foreign exchange, each event occurring with probability 0.5. Thus the outcomes in these sessions are random because the history of foreign inflation is.

3.3 Details of the Goods Markets

The experimenter also made the supply sides of both goods markets in every session. In essence, domestic goods and imports were both available in completely elastic supply at a price that was known to all. In sharp contrast with the last subsection, we will now concentrate on the *columns*, not the rows, of Table 1. Again, these columns are the experimental treatments. I begin by describing the control, the first column of Table 1. The upper left cell had two sessions that explored purchasing power parity; in this cell, the mark (domestic) price of German chocolate cake (domestic goods) was always one, and the franc (foreign) price of a loaf of French bread (foreign goods) was a treatment variable that assumed the following values: 0.25 francs, 0.5 francs, 1 franc, or 2 francs. This simple structure -- with perfectly elastic supply -- for the commodity markets induces an obvious derived demand for French francs, which was why the goods markets were designed to be so transparent

indeed. Every session of this experiment required that each subject spend all of his or her foreign exchange on imports at the end of each period. This aspect of the design was intended to circumvent foreign exchange hoarding and also to make the cash-in-advance constraint as stark as possible.

The middle left cell, whose session explored covered interest parity, had even simpler goods markets. Again the experimenter made the (perfectly elastic) supply sides of both commodity markets. The price of French bread was always one franc, and the price of German marks was also one mark. The point of the simple design was to induce a derived foreign exchange for French francs that depended upon the interest rate paid on franc deposits. This treatment examined covered interest parity because there was no risk to having an open position in francs, since the price of imports was always one franc.

The bottom left cell, whose sessions explored uncovered interest parity, set the price of German cake at one mark, and allowed the price of French bread to be one or two francs, depending upon the role of a die. Francs paid a (known) nominal interest rate, but holding an open position in foreign exchange (francs) was risky because the price of French bread could either be one or two, each with equal probability.

We now describe the cells in the second column of Table 1, the treatment having to do with not-traded goods. In these sessions, the subjects were paid according to:

$$u(b,c) = \alpha_b b + \alpha_c \sqrt{c} , \quad (1)$$

where b was the quantity of French bread consumed and c was that of German chocolate cake. The parameters α_b and α_c varied across sessions, and utility is obviously measured in dollars paid at the end of the experiment. Why is this equation

a good way to capture the notion of not-traded goods? In all the other treatments, the relevant utility function was:

$$u(b,c) = \alpha_b b + \alpha_c c, \quad (2)$$

where all the variables and parameters are analogous. Equation (2) shows that the marginal rate of substitution between bread and cake is constant, whereas equation (1) shows that this rate depends upon how many units of cake are consumed.

The simple form of (2) -- the payoff function in the treatments where all goods are traded -- has a very important implication. Let p_b be the franc price of a piece of bread, p_c be the mark price of a piece of chocolate cake, and e be the nominal exchange rate measured in German marks per French franc. Then the demand for German cake is zero if $ep_b / p_c < \alpha_b / \alpha_c$, and the demand for French bread is likewise zero if the opposite inequality holds. Thus if both goods are purchased in positive amounts -- which will be the case in equilibrium -- the only real exchange rate that would be observed is $ep_b / p_c = \alpha_b / \alpha_c$. So treatments where payoffs based on (2) cannot capture the effect of varying the real exchange rate that occurs when actual economies have forces that allow for smooth changes in the relative price of traded goods.

Let me hammer this point home. If nominal income is sufficiently high so that both goods are purchased, then the indirect utility function corresponding to (1) is:

$$v(p_b, p_c, m) = \alpha_b m / ep_b + (\alpha_c / 2)^2 ep_b / p_c, \quad (3)$$

where m is nominal income denominated in marks, and all the variables and parameters have the same interpretation as before. Equation (3) shows that the real exchange rate ep_b / p_c matters in these treatments. To my mind, this is a very simple formulation of the central idea in the seminal model of Salter (1959), where the real

exchange rate is described as the relative price of traded goods to not-traded goods. The non-linearity explicit in (1) captures the notion that domestic goods have a not-traded component. Consuming more domestically produced traded goods can only occur at an increasing opportunity cost, captured in these treatments by the simple artifice of declining marginal utility of consumption of German chocolate cake for agents endowed with German currency.

The only difference between the control and the treatment with not-traded goods is that subjects are paid according to (2) in the former and according to (1) in latter.¹¹ In particular, in sessions in the cell on top of the second column of Table 1, the domestic goods always cost one unit of domestic currency, but imports cost one of the following values: 0.25 francs, 0.5 francs, 1 franc, or 2 francs. Likewise, in the session in the middle cell of the middle column, German cake cost one mark and French bread cost one franc. In the sessions in the bottom cell of the middle column, domestic goods always cost one mark, but imports could cost one or two franc, with equal probability.

We can now complete this subsection by describing the sessions in the third column of Table 1, the treatment having a non-stationary environment. The good markets were identical to the control cases in every respect but one: the price of domestic goods (German chocolate cake) was 10 marks, then 20 marks, then 30 marks, and so on. Liquidity in domestic currency (German marks) was also increased *pari passu*, and, as always, German cake was in perfectly elastic supply. Since the franc price of French bread was just as in the control, this treatment predicts

¹¹ The subjects were given tables that specified exactly how much utility (in dollars paid out) they would receive for every possible number of pieces of German chocolate cake purchased. This table was just the numeric representation of the function $\alpha_c \sqrt{c}$, rounded to the nearest integer. The interested reader can look at the instructions to see the tables that were distributed.

that the German mark will depreciate in lock step with deterministic German “inflation.”

4. Experimental Procedures

The sessions were held in between October 2000 and June 2002 at The Ohio State University. The subjects were undergraduate and graduate students recruited by email from economics courses. The sessions typically involved six subjects, although there were four sessions with five subjects and one with only four.¹² The sessions lasted about ninety minutes, and they were conducted by hand, not by computer. The experiment is so simple that the experimenter and one assistant can run a session. The instructions for the nine different cells are available at <http://economics.sbs.ohio-state.edu/efisher/ppp/docs>.

Every session had two practice rounds and ten actual rounds. Thus the foreign exchange market cleared twelve times in the sessions in the first and second blocks of Table 1, and twenty-four times in all the other sessions. The subjects earned about \$25 on average, and the experiment earned a reputation for being “fun” perhaps because they involved simple tasks with which many subjects were already familiar.

Table 2 gives the actual parameters used in each of the cells. Perhaps this summary description of the parameters in each cell helps reinforce in the reader’s mind the two-way layout of the general experiment. In this table, the symbol $1 + i$ denotes the domestic gross nominal interest rate (which was always unity) and $1 + i^*$ denotes the counterpart foreign interest rate. The other symbols correspond to the discussion of the experimental design in the previous section.

¹² The session with four subjects occurred on 20 October 2000, and those with five subjects took place on 24 January 2002, on 15 February 2002, at 10:00 on 15 March 2002, and on 31 May 2002.

Table 2: Experimental Parameters				
		Macroeconomic Environment		
		Control	Not-Traded Goods	Non-Stationary Prices
<i>Element of Exchange Rate Theory</i>	Purchasing Power Parity	$(\alpha_b, \alpha_c) = (50, 5)$ $p_b \in \{1/4, 1/2, 1, 2\}$ $p_c = 1$	$(\alpha_b, \alpha_c) = (50, 5\sqrt{30})$ $p_b \in \{1/4, 1/2, 1, 2\}$ $p_c = 1$	$(\alpha_b, \alpha_c) = (50, 5)$ $p_b \in \{1/4, 1/2, 1, 2\}$ $p_c = 10, 20, \dots, 120$
	Covered Interest Parity	$(\alpha_b, \alpha_c) = (50, 5)$ $1 + i = 1$ $1 + i^* \in \{1, 2, 3\}$ $p_b = 1$ $p_c = 1$	$(\alpha_b, \alpha_c) = (25, 2.5\sqrt{30})$ $1 + i = 1$ $1 + i^* \in \{1, 2, 3\}$ $p_b = 1$ $p_c = 1$	$(\alpha_b, \alpha_c) = (25, 2.5)$ $1 + i = 1$ $1 + i^* \in \{1, 2, 3\}$ $p_b = 1$ $p_c = 10, 20, \dots, 120$
	Uncovered Interest Parity	$(\alpha_b, \alpha_c) = (50, 5)$ $1 + i = 1$ $1 + i^* \in \{1, 2, 3\}$ $p_b \in \{1, 2\}$ randomly $p_c = 1$	$(\alpha_b, \alpha_c) = (25, 2.5\sqrt{30})$ $1 + i = 1$ $1 + i^* \in \{1, 2, 3\}$ $p_b \in \{1, 2\}$ randomly $p_c = 1$	$(\alpha_b, \alpha_c) = (25, 2.5)$ $1 + i = 1$ $1 + i^* \in \{1, 2, 3\}$ $p_b \in \{1, 2\}$ randomly $p_c = 10, 20, \dots, 120$

The currencies were different kinds of play money. German marks were green play money, and French francs were red poker chips. The instructions made the two cash-in-advance constraints very explicit, and they also go into detail about how the call market works. After having read the instructions together with the experimenter aloud, all the subjects answered a series of questions to make sure that they understood the incentives inherent in each design.¹³ The answers to these questions were discussed publicly, and it was quite clear that the subjects understood how the call market worked. I repeatedly emphasized that the two low bidders get shut out of the foreign exchange market, but everyone else pays the same low price for foreign exchange.

¹³ The quizzes are part of the instructions for the sessions, and they can be found at <http://economics.sbs.ohio-state.edu/efisher/ppp>. In an important sense, these questions can be seen as helping to establish the common values in the third price auctions that constitute the foreign exchange markets. In no way did they offer any bidding strategy or indicate that it was a Nash equilibrium to bid this common value!

Very little record keeping was actually required of the subjects because they paid out their green currency and accumulated German chocolate cake and French bread as the session progressed. The experimenter gave each subject a large circle of orange construction paper; this big orange circle was called German chocolate cake, and using a special pen, I wrote on the big orange circle how many pieces of cake each subject bought in each period. This record was used for part of the reckoning of final payments. French bread was symbolized by oblong pieces of bright green construction paper, shaped like baguettes. The subjects just bought these loaves (and half loaves) and piled them up on their desks. At the end of the session, they were paid for each baguette and piece of chocolate cake they had bought.

I wrote out the history of prices for the French bread on the blackboard, and I also wrote out the prices (namely, the green currency price of a red poker chip) at which the foreign exchange market cleared in each period or sub-period. For every transaction, a subject spent actual currency to buy goods or foreign exchange. The experimenter was always on the opposite side of each transaction, and the subjects were always endowed only with green currency.

The subjects submitted their bids on small chits of paper that were collected and put in rank order. Using a simple Excel file, I was able almost immediately to announce the market-clearing price (the second lowest bid) for foreign exchange.¹⁴ As is usual with a call market, the infra-marginal bids were not made public information. In the case of ties among the marginal bids, a public randomization device (a die) decided who could buy foreign exchange. Bids for foreign exchange were for one red poker chip.

¹⁴ These files also include all the data. The interested reader can find them at <http://economics.sbs.ohio-state.edu/efisher/ppp/data>.

For the treatments where French francs bore interest, the nominal rate of return was announced before the beginning of sub-period A. Subjects who had purchased red poker chips in sub-period A were given extra red poker chips in time to go shopping at the end of sub-period B. In the treatment exploring uncovered interest parity, the price of French bread was random, determined according to a public toss of a die. If 1, 2, or 3 came up, the price of French bread was one franc. If 4, 5, or 6 came up, the price of French bread was two francs.

The procedures were really that simple, and I could run them with one assistant.

5. Equilibrium in the Foreign Exchange Market

The no-surplus Nash equilibrium forms the theoretical predictions for the model. Each period in each session can be described as a game in normal form. The subjects are bidding for foreign exchange in a common value auction. Let there be H_i subjects in the i -th session, and write $b_i^h(t)$ as the bid of subject $h \in \{1, \dots, H_i\}$ at time t in that session. Each foreign exchange market is a third-price auction in which the highest $H_i - 2$ bidders buy foreign exchange at the marginal excluded bid.

Let the price at which no subject derives any surplus from buying foreign exchange in period t of the i -th session be $y_i(t)$. In the diagrams below, $y_i(t)$ will be shown as the theoretical prediction for the exchange rate. I claim it is a (symmetric) Nash equilibrium for every person to bid this price. Assume that $H_i - 1$ subjects are all bidding this value.¹⁵ If $b_i^h(t) > y_i(t)$, then subject h definitely receives a unit of foreign exchange, but he has no effect on the price and still derives no surplus from it in the common value auction. If $b_i^h(t) < y_i(t)$, then subject h bids

too low, cannot purchase foreign exchange, and again derives no surplus from that market. We have shown that bidding $b_i^h(t) = y_i(t)$ for all subjects $h \in \{1, \dots, H_i\}$ is a Nash equilibrium in this game.¹⁶

Unfortunately, this is not the only Nash equilibrium in the one-shot game, and the other equilibria are not observationally equivalent when analyzing either individual data like $b_i^h(t)$ or aggregate data like the actual market price, which we will call $x_i(t)$. Assume that $H_i - 2$ subjects bid the theoretical no-surplus prediction $y_i(t)$, and for argument's sake posit that $b_i^h(t) > y_i(t)$. Let $j \neq h$ be some other subject. Then if $b_i^j(t) \geq b_i^h(t) > y_i(t)$, both subjects h and j definitely purchase foreign exchange but neither derives any surplus because the marginal excluded bid is still $y_i(t)$. Likewise, none of the other $H_i - 2$ subjects has any incentive to change his bid, since a deviation in either direction will not change the market price and create any possibility of strictly positive surplus.¹⁷ Thus we have shown that the Nash equilibrium in this game is not unique, and that different data on individual bids give rise to the same market outcome.

It is even more troublesome that there are some Nash equilibria that do not give rise to the theoretical market price $y_i(t)$. Until now all the arguments have been based upon the no-surplus principle. Assume now that $H_i - 2$ subjects bid $\bar{x}_i(t)$, the maximum allowed in period t of the i -th session. (This maximum was determined by the German mark liquidity that was provided in the different treatments.) Assume that subjects $h \neq j$ both bid $b_i^j(t) = b_i^h(t) = 0$. Then the market price for foreign

¹⁵ Please recall that each person purchases foreign exchange with probability less than unity if everyone bids the same value. Still, expected surplus is zero for any buyer or for any person who bids too low.

¹⁶ Indeed, if we think of each session as a repeated game, it is a sub-game perfect Nash equilibrium for every trader to bid the predicted market price every time since there are only finitely many periods in any treatment.

exchange is $x_i(t) = 0$, and all the subjects bidding $\bar{x}_i(t)$ get plenty of surplus.¹⁸ I claim that this configuration of bids is a Nash equilibrium in the one-shot game. Here is why. None of the subjects bidding $\bar{x}_i(t)$ has any incentive to change his bid, since each of them gets foreign exchange with probability one. Likewise, neither h nor j has any incentive to bid higher because neither has any chance of buying foreign exchange for any bid strictly less than $\bar{x}_i(t)$. Further, bidding the maximum actually raises the market price to that level, causing non-positive expected surplus since $\bar{x}_i(t) \geq y_i(t)$.¹⁹

This section has established two important points. First, the symmetric no-surplus equilibrium in every period is a Nash equilibrium in the finitely repeated game that each session constitutes. This equilibrium is the natural basis for the theoretical predictions for each session. Second, there are other Nash equilibria, some of which have the same no-surplus market prices and some of which do not. All of these asymmetric equilibria exhibit dispersion among the individual bids that the subjects submit. With these facts firmly in mind, we can now turn our attention to the data.

6. Experimental Results

The results are best summarized by graphs. Figures 1 through 5 show the predictions and the actual data for all five sessions in the control treatment, Figures 6 through 10 show the analogous five sessions for the treatment with not-traded goods, and Figures 11 through 15 show those for the treatment with non-stationary domestic

¹⁷ I am assuming that there are at least 5 subjects in this session.

¹⁸ Also, the experimenter pays out a lot more grant money than he expected, especially if the subjects focus on this equilibrium or one of its variants as a focal point for the sub-game perfect equilibrium in the finitely repeated game!

¹⁹ The liquidity constraint in domestic currency was binding no more than four times in all the periods of all these sessions. Thus bidding the maximum in such a situation would almost always entail strictly negative expected profits.

prices. In all these figures, the predicted exchange rate is quite near the actual data. In several sessions, there are deviations from the theoretical predictions, but the five sessions that constitute the control treatment have remarkably accurate theoretical predictions. In an important sense, this finding is serendipitous because it allows me to identify *which* elements of exchange rate theory fare (relatively) poorly in laboratory analysis. The sessions from the treatments with not-traded goods and with non-stationary prices do show greater deviations from the theoretical predictions than those in the control group. And the very structure of the experiment allows me to identify in a controlled environment the causes of deviations from the theoretical predictions that form some of the fundamental elements of exchange rate theory.

Let us consider some general observations about these fifteen sessions. To my mind, three points are salient. First, the purchasing power parity, covered interest parity, and uncovered interest parity -- three elements of exchange rate theory -- fare very well in these sessions. The data from this experiment conform much better to the theoretical predictions than is the norm in the econometric analysis of exchange rate models that use field data. Second, the control sessions -- those described in the first column of Table 1 and shown in Figures 1 through 5 here -- exhibit remarkably accurate theoretical predictions.²⁰ Thus risk-neutral subjects can perform the cognitively simple task of giving up domestic currency to purchase foreign exchange

²⁰ Why do these control sessions fare so well, when Noussair, Plott, and Riezman (1997) found weak or perhaps mixed evidence in support of their model of exchange rate determination? My experiment has a much simpler design that gets at the essence of some elements of exchange rate theory. In particular, these subjects can focus entirely on bidding for and thus buying foreign exchange since the experimenter makes the supply of that market. Likewise, no subject really has to worry about earning domestic currency by selling goods, since the experimenter also makes that supply side of both goods markets. A simple design is, of course, the hallmark of a good experiment. I do not mean to say that the conclusions of Noussair, Plott, and Riezman are erroneous; instead the sessions that constitute my control seem to indicate that purchasing power parity failed in previous experiments because of the interaction of the goods markets and the market for foreign exchange. It is a commonplace to observe that the link between commodity markets and asset markets is very subtle in economic theory, let alone in the laboratory. The contrast between my results and those of Noussair, Riezman, and Plott reinforces this observation.

if the controlled economic environment is straightforward. Third, both not-traded goods and non-stationary domestic prices cause deviations from the theoretical predictions, especially (although not exclusively) during periods of large depreciations. Compare the sessions in the treatment with not-traded goods (depicted in Figures 6 through 10) with their analogs in the treatment with non-stationary domestic prices (depicted in Figure 11 through 15). This comparison permits me to highlight perhaps the most important contribution of this paper to the empirical literature on exchange rate theory. *The biggest deviations from the theoretical predictions seem to occur because of the non-stationary environment.*

I shall reinforce this observation with the appropriate statistical analysis. The relevant data are given in Tables 3 and 4. Again, let $x_i(t)$ be the data from period t for the i -th session, and recall that $y_i(t)$ is the analogous theoretical prediction. I analyze two natural measures of the goodness of fit of the theory. The first is:

$$z_i = \sqrt{\sum_{t=1}^{T_i} (x_i(t) - y_i(t))^2 / T_i}, \quad (3)$$

where $T_i \in \{10,20\}$ is the number of periods or sub-periods in the i -th session. Thus (3) describes the root mean squared error for a given session. These goodness-of-fit statistics are given in Table 3, and recall that they place weight on large outliers.

Table 3: Goodness of Fit Statistics, Root Mean Squared Errors				
		<i>Macroeconomic Environment</i>		
		Control	Not-Traded Goods	Non-Stationary Prices
<i>Element of Exchange Rate Theory</i>	Purchasing Power Parity	3.12	5.62	63.45
	Covered Interest Parity	4.75	1.18	41.07
	Uncovered Interest Parity	1.14	2.39	49.07
	Uncovered Interest Parity	1.77	5.78	21.24
		1.19	3.66	54.30

The second natural goodness-of-fit statistic is:

$$\tilde{z}_i = (1/T_i) \sum_{t=1}^{T_i} \left| \frac{x_i(t) - y_i(t)}{y_i(t)} \right| \quad (4)$$

where the variables are analogous. This measure is the mean absolute percentage errors, and it normalizes each observation by the theoretical prediction, perhaps controlling for the fact that large absolute errors occur for extreme values of the theoretical predictions, especially in models with non-stationary outcomes. These statistics are given in Table 4.

Table 4: Goodness of Fit Statistics, Mean Absolute Percentage Errors				
		<i>Macroeconomic Environment</i>		
		Control	Not-Traded Goods	Non-Stationary Prices
<i>Element of Exchange Rate Theory</i>	Purchasing Power Parity	13%	28%	33%
	Covered Interest Parity	8%	8%	26%
	Uncovered Interest Parity	2%	7%	26%
	Uncovered Interest Parity	15%	38%	19%
		5%	22%	67%

6.1 Statistical Analysis of the Treatment Effects

The experiment was designed specifically to give a balanced panel of data. Each goodness-of-fit observation is of course independent because the sessions were done with different samples of a (fairly homogenous) population of subjects. It is worth mentioning that there may be some ambiguity about the exact definition of a bloc and a treatment in any balanced design. I think of this panel as having five blocs and three treatments, although one can use an alternative interpretation of three blocs and five treatments. My preferred interpretation is that the treatments are the three different macroeconomic environments: the control, the existence of not-traded goods, and the existence of non-stationary (domestic) prices. In this case, a test for treatment effects is a test for whether the macroeconomic environment matters for all

three elements of exchange rate theory *in toto*. Under the alternative specification, there are three blocs and five treatments, and a test for treatment effects has the interpretation that there is something special about one of the three elements of exchange rate theory: purchasing power parity, covered interest parity, or uncovered interest parity. The primary conclusion of this subsection is that there are treatment effects in these data. A secondary conclusion is that there is nothing special about any one of the three elements of exchange rate theory that I explore.

The natural non-parametric test for treatment effects is Friedman's (1937) rank sum test. It is based upon the ranks of the outcomes in each of the five rows. Under the null hypothesis that there are no treatment effects, the ranks in each row should be random, and the column sums of the ranks ought to show no clear pattern. The alternative hypothesis is that at least one treatment has some kind of effect. Now look closely at the data in Tables 3 and 4. The observations in the second and third columns have large ranks; thus there is very strong evidence in favor of treatment effects. Indeed, using the data from either Table 3 or Table 4, one can reject the null hypothesis of no treatment effect for a test of size 1%.²¹

Consider the transpose of the data in Tables 3 and 4. Using the transposed data, I calculated the Friedman rank sum statistic. The p-value of the statistic for the transposed data from Table 3 is .559, and that for the transposed data in Table 4 is .291. Thus there is little evidence that there is anything special about any of the three elements of exchange rate theory that I am analyzing. Henceforth, I will always consider the macroeconomic environment as the treatment variable and explain in general why all three elements of the theory seem to fare poorly in one

²¹ In both cases, the Friedman statistic is 8.4, and it has a p-value of .008. There were no ties; the data shown in Table 4 have been rounded for ease of exposition. The exact distribution of this statistic is given in Hollander and Wolfe (1973), Table A.15. It has an asymptotic chi-squared distribution whose degrees of freedom are one less than the number of treatments.

macroeconomic environment or another. Still, the secondary conclusion of this subsection is that purchasing power parity, covered interest parity, and uncovered interest parity fare roughly equally well in these data, in each of the three macroeconomic environments.

6.2 Which Effect is Stronger, Not-Traded Goods or Non-Stationary Prices?

The statistics in the third column of Table 3 are much larger than those in the second column. That table seems to present *prima facie* evidence that the largest cause for departures from the theory has to do with non-stationary domestic prices. Still, there is a nagging problem with the interpretation of the root mean squared error in a non-stationary environment: that statistic puts large weight on extreme outliers. This is the main reason that I chose also to analyze the mean absolute deviations in Table 4.

The natural non-parametric test for ordered treatment effects is Page's (1963) test. The formal model for the data (see Hollander and Wolfe (1973), p.139) is:

$$X_{ij} = \mu + \beta_i + \tau_j + e_{ij}, \quad i = 1, \dots, 5, \quad j = 1, 2, 3, \quad (5)$$

where μ is the overall mean, β_i is a nuisance parameter that measures the block effect, τ_j is the treatment effect and the e_{ij} are mutually independent. The null hypothesis for both Friedman and Page's test is:

$$H_0 : \tau_1 = \tau_2 = \tau_3,$$

but the alternative hypothesis for Page's test is:

$$H_0 : \tau_1 \leq \tau_2 \leq \tau_3.$$

Page's test is based upon the statistic $S = \sum_j jR_j$, where R_j is the sum of the ranks in column j . Using the data from Table 3 and from Table 4, I calculated the Page statistic, and it had a value of 69 in both cases. This is large enough to reject the null

hypothesis in favor of the alternative for a test of size 1%. Hence, there is strong evidence in these data that there are treatment effects and also that the effect of non-stationary prices is at least as important as that of not-traded goods.

6.3 Why is the Effect of Non-Stationary Prices Stronger?

Why does the treatment with non-stationary prices cause relatively large deviations from the theoretical predictions? Perhaps it is a more cognitively demanding task to purchase foreign exchange in a non-stationary environment. Even though payoffs in the treatment with not-traded goods are non-linear, it is quite plausible that the subjects find it less complex to purchase foreign exchange in an environment where neither domestic prices nor foreign prices have drift. I examined that conjecture by analyzing the bids of all the subjects in each session. Let $b_i^h(t)$ be the bid of subject h in period t of the i -th session, write H_i as the number of subjects in that session.²² Again, let $y_i(t)$ be the theoretical prediction for market price in period t of the i -th session. Then I used

$$b_i = (1/T_i)(1/H_i) \sum_{t=1}^{T_i} \sum_{h=1}^{H_i} \left| \frac{b_i^h(t) - y_i(t)}{y_i(t)} \right| \quad (6)$$

as the basis for my tests. The statistic in (6) is the mean -- across traders and periods -- absolute percentage error of the bids from the theoretical market price. It controls for the fact that the theoretical predictions in the non-stationary environment have a deterministic drift.

Section 5 established that there are several Nash equilibria in the one-shot game. It is now time to examine the individual bids more closely. Table 5 gives the statistic in equation (6). The data in Table 5 are similar to those in Table 4, but they

²² This kind of analysis shows an obvious advantage of experimental economics. Imagine how hard it would be for the econometrician to obtain data on all the (marginal and infra-marginal) bids in an actual foreign exchange market!

are more dispersed. This increased dispersion indicates that the treatment effect on the individual bidders in general is more pronounced than that on the marginal bidders.

		<i>Macroeconomic Environment</i>		
		Control	Not-Traded Goods	Non-Stationary Prices
<i>Element of Exchange Rate Theory</i>	Purchasing Power Parity	37%	48%	44%
	Covered Interest Parity	14%	54%	78%
	Uncovered Interest Parity	9%	15%	60%
	Uncovered Interest Parity	26%	58%	65%
		53%	46%	200%

I explored the conjecture that the subjects find it a more cognitively challenging task to make bids in the non-stationary environment by again using Page's (1963) test. The reader is reminded (gently) that the null hypothesis is the lack of a treatment effect and the alternative is that the treatment effects are monotonic, with the non-stationary one having the strongest effect and the control the weakest one. The value of the Page statistic for the data from Table 5 is now 68, and I can reject the null for a test of size 1%.

I feel fairly confident that the subjects making the infra-marginal bids are somewhat "confused" because the basic theoretical predictions are accurate enough, as one can see from the Figures 11 through 15 or from the data in Table 4. These data indicate marginal bidders are bidding much closer to the predicted price than the average bidder is, especially in the non-stationary treatments. Since the predicted price is based upon the no surplus Nash equilibrium, it does not seem to be the case that the marginal bidders are acting strategically in the treatments where the infra-marginal bidders are making widely divergent bids. Indeed, in almost every session,

there was very little dispersion in the payments that I actually made to the different subjects, indicating those few if any subjects were earning large infra-marginal rents.

Before concluding this sub-section, I would like to interject a note of caution.²³ I have argued--somewhat loosely--that the non-stationary prices cause much greater deviations from the theoretical predictions than do not-traded goods. But I am in the same position as a medical researcher who has given doses of two different drugs to members of three related species. I know that the mortality rates are much higher for the second drug (non-stationary prices) than for the placebo (the control group) or the first drug (not-traded goods). I know that these effects hold for each of the three different but related species in my experiment. But how does one control for the dosages? In particular, perhaps I have administered small doses of the first drug and a gigantic dose of the second. So I must admit that a more measured conclusion is that the magnitude of non-stationary domestic prices in this particular experiment causes greater deviations from the theoretical predictions than does the magnitude of the non-linearity that captures the macroeconomics of not-traded goods. Whether this effect is robust to different doses in future experiments is obviously an open area of research.

7. Conclusion

This paper has used the laboratory to show that elements of exchange rate theory perform fairly well in simple environments. One important conclusion is that there is nothing special about purchasing power parity, covered interest parity, or uncovered interest parity. A simple enough environment -- like the control treatment -- shows that each of these building blocks of any model of exchange rates holds up well in the laboratory. So an important conclusion of this work is that Hypothesis 1,

²³ This paragraph was directly inspired by a conversation with Charles B. Grant.

2, 3, and 4 are true: absolute purchasing power parity, relative purchasing power parity, covered interest parity, and uncovered interest parity all perform remarkably well in a simple cash-in-advance model in the laboratory. I think Gustav Cassel or Irving Fisher would be reassured by the sessions in the control treatment, but then again there is nothing really dramatic about finding that these subjects don't suffer from money illusion, understand simple concepts of interest rates, and are essentially risk-neutral for small bets. On the other hand, the control treatment indicates the somewhat negative findings of Noussair, Riezman and Plott (1997) or Arifovic (1996) have perhaps much to do with the interaction between the goods and asset markets, not with elements of exchange rate theory *per se*.

The primary conclusion I draw from this work is that people seem to find it more difficult to derive a common value for foreign exchange in a simple non-stationary environment than they do in a simple stationary non-linear one. A tentative (and admittedly quite speculative) conclusion for the policy maker might be that the modern period of floating rates has seen such wide swings of real exchange rates because countries have had different inflation experiences during the last quarter of the twentieth century. I think that further study into the cognitive psychology of foreign exchange traders might be very instructive. Survey research like that reported by Frankel and Froot (1987) is quite useful, but economic experiments are too.

I would like to conclude with an exhortation for more experimental research in macroeconomics and international economics. The hallmark of a science is that the empirical implications of its theories can be replicated. The experiment that I describe here is just the first step into what I hope will become a much broader research program. I have shown that the elements of exchange rate theories are vindicated in the laboratory (as well they should be in a simple enough design). But I

have also used a controlled environment to indicate perhaps that the broad empirical failure of many exchange rate models may have to do with the disparate secular inflation rates that the major industrial have experienced. Perhaps the current convergence of inflation rates may herald a coming decade of exchange rate data in which the models and the empirical research will not be at such sixes and sevens. Of course, it may just be my own hubris to hope that this vein of (highly stylized) research has shed some light on the actual functioning of the international monetary system. But I do hope that I have piqued some reader's curiosity enough so that you might contribute to this nascent literature in the not distant future.

References

- Arifovic, J. "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies," *Journal of Political Economy* 104 (1996), 510-41.
- Cassel, G. "The Present Situation of the Foreign Exchanges," *Economic Journal* 26 (1916), 62-5.
- Engel, Charles. "Accounting for U.S. Real Exchange Rate Changes," *Journal of Political Economy* 107 (1999), 507-38.
- Faust, Jon R., John H. Rogers, and Jonathan H. Wright. "Exchange Rate Forecasting: The Errors We've Really Made," International Finance Discussion Paper 714, 2001.
- Fisher, Eric O'N. "Purchasing Power Parity and Interest Parity in the Laboratory," *Australian Economic Papers* 40 (2001), 586-602.
- Fisher, Eric O'N. and F. S. Kelly. "Experimental Foreign Exchange Markets," *Pacific Economic Review* 5 (2000), 368-88.
- Fisher, Eric O'N. and Joon Y. Park. "Testing Purchasing Power Parity under the Null Hypothesis of Co-integration," *Economic Journal* 101 (1991), 1476-84.
- Flood, Robert P. and Peter M. Garber. "Market Fundamentals versus Price-Level Bubbles: The First Tests," *Journal of Political Economy* 88 (1980), 745-70.
- Frankel, Jeffrey A. "On the Mark: A Theory of Floating Exchange Rates Based on Real Interest Differentials," *American Economic Review* 69 (1979), 610-22

Frankel, Jeffrey A. and K. A. Froot. "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *American Economic Review* 77 (1987), 133-53.

Friedman, Milton. "The Use of Ranks to Avoid the Assumption of Normality Implicit in the Analysis of Variance," *Journal of the American Statistical Association* 32 (1937), 675-701.

Hollander, Myles and Douglas A. Wolfe. *Nonparametric Statistical Methods*. John Wiley: New York, 1973.

Mark, Nelson C. "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *American Economic Review* 85, 201-18.

Meese, Richard A. and K. Rogoff. "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics* 14 (1983), 3-24.

Neeley, Christopher J. and L. Sarno. "How Well do Monetary Fundamentals Forecast Exchange Rates?" unpublished manuscript, 2002.

Noussair, Charles, C. Plott, and R. Riezman. "The Principle of Exchange Rate Determination in an International Finance Experiment," *Journal of Political Economy* 105 (1997), 822-61.

Page, E. B. "Ordered Hypotheses for Multiple Treatments: A Significance Test for Linear Ranks." *Journal of the American Statistical Association* 58 (1963), 216-230.

Salter, W. "Internal and External Balance: The Role of Price and Expenditure Effects." *Economic Record* 36 (1959), 51-66.




Figure 1
Purchasing Power Parity, 5 October 00




Figure 2
Purchasing Power Parity, 10 October 2000




Figure 3
Covered Interest Parity, 13 October 2000




Figure 4
Uncovered Interest Parity, 20 October 2000




Figure 5
Uncovered Interest Parity, 24 January 2002




Figure 6
Not Traded Goods Purchasing Power Parity, 8 February 2002




Figure 7
Not Traded Goods Purchasing Power Parity, 15 February 2002




Figure 8
Not Traded Goods Covered Interest Parity, 8 March 2002




Figure 9
Not Traded Goods Uncovered Interest Parity, 30 May 2002




Figure 10
Not Traded Goods, Uncovered Interest Parity, 7 June 2002




Figure 11
Non-Stationary Purchasing Power Parity, 1 February 2002




Figure 12
Non-Stationary Purchasing Power Parity, 1 March 2002




Figure 13
Non-Stationary Covered Interest Parity, 22 February 2002




Figure 14
Non-Stationary Uncovered Interest Parity, 10:00 15 March 2002




Figure 15
Non-Stationary Uncovered Interest Parity, 14:00 15 March 2002