Macroeconomic Dynamics, **5**, 2001, 1–31. Printed in the United States of America.

ARTICLES

CONVERGENCE IN MONETARY INFLATION MODELS WITH HETEROGENEOUS LEARNING RULES

GEORGE W. EVANS
University of Oregon

SEPPO HONKAPOHJA
University of Helsinki

RAMON MARIMON
European University Institute
and
Universitat Pompeu Fabra

Inflation and the monetary financing of deficits are analyzed in a model in which the deficit is constrained to be less than a given fraction of a measure of aggregate market activity. Depending on parameter values, the model can have multiple steady states. Under adaptive learning with heterogeneous learning rules, there is convergence to a subset of these steady states. In some cases, a high-inflation constrained steady state will emerge. However, with a sufficiently tight fiscal constraint, the low-inflation steady state is globally stable. We provide experimental evidence in support of our theoretical results.

Keywords: Inflation, Heterogeneous Learning, Heterogeneous Expectations, Fiscal Constraints, Multiple Equilibria

1. INTRODUCTION

Recently, economists have been interested in the role of constitutional limitations on public-sector deficits and debt. An example is provided by the fiscal "convergence criteria" in the Maastricht plan and the Stability Pact for the European Monetary Union, which put an upper limit on the deficit relative to GDP. Most

The referees provided useful comments on an earlier version. We are also grateful for comments received at numerous seminars, workshops, and conferences. The experiments reported here were performed at the experimental lab of the Universitat Pompeu Fabra (LeeX) in June 1994, using the software developed by R. Marimon and S. Sunder. We would like to thank Carles Perarnau for his assistance, the Spanish Ministry of Education for the financial support (through a DGICYT grant), and the UPF undergraduates who participated in the experiments. We acknowledge financial support from the SPES program of the EU, the National Science Foundation, the Yrjö Jahnsson Foundation, and the Academy of Finland. Address correspondence to: Professor George W. Evans, Department of Economics, 1285 University of Oregon, Eugene, OR 97403-1285, USA; e-mail: gevans@oregon.uoregon.edu.

 1365-1005/01 $9.50

of the literature has focused on either the incentive aspects of such constraints or their consequences for policies to stabilize the business cycle. In this paper, we take a different track and consider the implications of a credible constraint on public deficits for monetary stability.

The analysis is an extension of the standard deficit-finance monetary inflation model in which expectations are modeled as an outcome of a learning process.[1] Adaptive learning dynamics have attractive properties and we make two contributions to the literature on learning. First, our analysis shows that it is important to take the process of expectation formation into account in the design of economic policies. Though the introduction of a fiscal constraint aggravates the indeterminacy of rational expectations equilibria, the constraint has a stabilizing role for inflation under learning and the resulting policy implications can be striking. Second, we obtain stability results for an economy with heterogeneous learning rules. We also provide experimental support for the theoretical analysis of the model under learning.

We begin by reviewing the standard deficit-finance monetary model. The starting point is the government budget constraint, which, in per-capita terms, takes the form

$$h_t - h_{t-1} = p_t d_t, \tag{1}$$

where h_t, p_t, and d_t denote per-capita nominal money balances, the price level, and the real deficit, respectively. It is usually assumed that the demand for real balances, $m_t = h_t/p_t$, depends only on the expected return on money holdings, $m_t = m(\pi_{t+1}^e)$, where $\pi_{t+1} = p_{t+1}/p_t$ is the rate of inflation and π_{t+1}^e denotes expected inflation. Rearranging the budget constraint as $m_t = m_{t-1}/\pi_t + d_t$ we can solve for π_t. This yields the monetary equilibrium equation

$$\pi_t = \frac{m\left(\pi_t^e\right)}{m\left(\pi_{t+1}^e\right) - d_t},$$

to which is traditionally added the perfect-foresight hypothesis $\pi_t^e = \pi_t$.

The earlier literature [see, e.g., Marcet and Sargent (1989a) and Sargent and Wallace (1987)] has usually postulated a given real public deficit $d_t = d$ financed by printing money. For simplicity, it is typically assumed that money demand is linear. In the standard formulation, provided d is not too large, there are two perfect-foresight steady states, π^A and π^B, with $\pi^B > \pi^A$. (If d is sufficiently large, the model has no perfect-foresight solutions other than autarky and worthless money.) In addition, there is a continuum of perfect-foresight paths converging to π^B. This observation has been the basis of a large literature interpreting empirical episodes of high or hyper-inflation as convergence to π^B [see, e.g., Fischer (1984), Bruno (1989), and Sargent and Wallace (1987)].

There are, however, difficulties with this view. First, π^B can be shown to have paradoxical comparative static properties. For example, an increase in d leads to a reduction in π^B. In contrast, π^A has conventional comparative statics. The second difficulty concerns the behavior of the model under adaptive learning. For this model, it can be shown that π^B is unstable under most adaptive learning rules [see, e.g., Marcet and Sargent (1989a)]. They show that under their learning scheme there are two possibilities. Inflation either converges to π^A or else becomes so high that the deficit cannot be financed at any finite price level, in effect resulting in currency collapse. (For large values of d, only the latter outcome arises.)[2]

A principal objective of this paper is to use this model to analyze the implications of an additional fiscal constraint on the size of the deficit relative to a measure of aggregate market activity. (We prefer to think of this as GDP, though different interpretations of this aggregate variable are possible, as discussed later.) We show that this rule can create a new steady state, π^C, with a finite inflation rate, in the constrained regime in which the government is not able to fully achieve its target deficit. When it exists, π^C is locally stable under natural adaptive learning rules. This feature leads to the following economic results.

If the target deficit is large, the economy converges to the new steady state π^C of the constrained regime. If the target deficit is not too large and the fiscal constraint is not very strict, the steady states π^A and π^B also exist in our model, and their local stability properties under learning are as in the standard model. In our model, there is thus the possibility of two locally stable steady states π^C and π^A under learning. (Note that, in the standard model, a large d leads to eventual currency collapse, and this can also arise at lower d.) Another possibility is that the constraint is set sufficiently tightly that π^B no longer exists (because it does not satisfy the fiscal constraint) but π^A does. Then, π^A is the only stable steady state, and it is globally stable. Comparing this to the preceding case, it follows that a tightening of the deficit constraint can move the economy from high inflation to a stable steady state with low-inflation.[3] If the fiscal contraint is made even stricter, the unconstrained low-inflation steady state is converted into a constrained type.

These results indicate the potential importance of fiscal constraints for monetary stability. An appropriate constitutional limitation on the relative size of the deficit can ensure that the economy has a unique globally stable steady state at a low inflation rate. We show that, in all cases, there is global convergence under learning to the set of locally stable steady states in which money has value.

In showing these formal results, we significantly extend the literature by allowing heterogeneous learning rules. Though developed in a concrete monetary model, our methodology has general validity. An additional reason for treating heterogeneous expectations is its clear empirical relevance. Indeed, experimental evidence shows that even when agents are, by design, homogeneous in terms of utility and budget sets, their expectations differ substantially.

In the last section of the paper we provide new experimental results. The experimental design is very close to the specific model studied in this paper. Agents have homogeneous (induced) preferences and endowments. Nevertheless, the subjects do exhibit heterogeneity in their forecasts. The experimental data provide support for various aspects of our theoretical analysis. In particular, the economic results on the stability properties of the steady states described in the preceding paragraph are strikingly reflected in the experimental results.

2. MONETARY MODEL

As is common in the literature, we develop a monetary overlapping generations (OLG) model of deficit finance in which we assume that the deficit is financed through the inflation tax.[4] We add the feature that the actual deficit cannot exceed a certain level, as specified later. We can think of the choice of such a fraction as a more fundamental decision taken by society—for example, in the constitution—and that it constrains the monetary authority.

More formally, at any point in time, there is a continuum of agents located on the unit interval [0, 1] with the Lebesgue measure denoted by μ. In period t, the demand for money of agent $k \in [0, 1]$ is given by

$$h_{kt} = \max\left\{ p_t\left(a - b\pi_{t+1}^{ke}\right),\ 0\right\}, \tag{2}$$

where π_{t+1}^{ke} denotes k's expectation, at the beginning of period t, about the rate of inflation between periods t and $t+1$. It is easy to see that (2) can be derived from an OLG model in which agent k has a two-period endowment of a unique perishable good $(\omega^1,\ \omega^2)$, where $\omega^1 > \omega^2 > 0$, with preferences over consumption represented by $u(c^1,\ c^2) = \ln(c^1) + \beta \ln(c^2)$. Here, the superscript denotes the period in the agent's life. To focus on the role of expectations, we assume that agents have identical preferences and endowments, but can differ in expectations.[5]

An agent k of generation t, $t \geq 1$, solves the problem

$$\begin{aligned} \max \quad & \ln c_t^1 + \beta \ln c_t^2 \\ \text{s.t.} \quad & c_t^1 - \omega^1 + \pi_{t+1}^{ke}\left(c_t^2 - \omega^2\right) \leq 0. \end{aligned}$$

If $\omega^1 - \omega^2$ is large enough, k's supply in the first period of his life is

$$s_{t,k} = \max\left\{ \left(\beta\omega^1 - \pi_{t+1}^{ke}\omega^2\right)/(1+\beta),\ 0\right\}.$$

That is, in our general formulation, $a = \beta\omega^1/(1+\beta)$ and $b = \omega^2/(1+\beta)$. Therefore, the average savings by the young in period t are given by

$$s_t = \int_{[0,1]} s_{t,k}\mu(dk) = \int_{[0,1]} \max\left\{a - b\pi_{t+1}^{ke},\ 0\right\}\mu(dk). \tag{3}$$

Although the preceding presentation in terms of endowments of goods is the usual one, it is convenient to give an alternative interpretation in terms of production. On this interpretation, ω^1 and ω^2 represent endowments of labor time in the two periods. One unit of time produces one unit of the good, and the household can either directly consume the goods produced or sell part of them on the market. Under our assumptions, only a portion of the goods produced by the young in period t enters the market, specifically $\omega^1 - c_t^1$ per young household. It is natural to identify this total amount of traded goods as GDP, following the usual accounting practice that nonmarket production is excluded from measured GDP. Thus, in this interpretation, GDP corresponds to the aggregate savings s_t of the young.[6]

If agents have the same beliefs (i.e., $\pi^{ke}_{t+1} = \pi^{e}_{t+1}$), then (3) can be simplified. The *rational-expectations hypothesis* imposes this form of homogeneity of beliefs by assuming that $\pi^{ke}_{t+1} = \pi_{t+1}$. In this section, we describe the rational-expectations steady states and, in the rest of the paper, we assume that the possibly heterogeneous expectations are formed through a process of adaptive learning. Although throughout we assume homogeneous preferences and endowments, our analysis can be easily extended to the case of a finite number of agents' types with respect to preferences and endowments.

The aggregate savings function (3) provides a complete characterization of the behavior of the private sector in this economy. As in the standard inflation model the government finances a real deficit by increasing the per-capita supply of money.[7] We modify the usual model in two respects. First, and most fundamentally, it is assumed that there is some limitation on the permissible deficits. Apart from transitory shocks, the actual deficit is required to be no greater than a certain fraction λ of GDP. We therefore assume that the government has a target level of the deficit, d_t, which is fully financed by seignorage, provided the preceding restriction is not violated. If instead the restriction is binding, government expenditures are reduced as necessary to meet the constraint on deficits. As discussed earlier, in this model, measured GDP and aggregate savings are equal. Furthermore, savings are equal to real balances, so that our constraint could be interpreted in any of three ways. The key feature is the imposition of a constraint that real deficits cannot exceed a fraction of an endogenous variable associated with aggregate market activity.[8]

The second modification is that we introduce transitory random shocks to government expenditure. Let g_t be an i.i.d. random variable with a small compact support and $Eg_t = 1$. If h_t denotes the per-capita money supply at the end of period t, the supply of money evolves according to

$$h_t = g_t h_{t-1} + \min\{p_t d_t,\, p_t \lambda s_t\}.$$

If $g_t = 1$ and $d_t < \lambda s_t$, this equation reduces to equation (1) in the standard model. When $\lambda s_t < d_t$, the postulated restriction forces a reduction in the actual deficit below the target level d_t through reductions in government spending. Values of $g_t \neq 1$ allow in a convenient way for additional random shocks to spending financed by money creation.

The preceding equation can be written in terms of real balances as

$$m_t = g_t m_{t-1}/\pi_t + \min\{d_t, \lambda s_t\}, \tag{4}$$

where $m_t = h_t/p_t$ is the per-capita money supply in real terms. In this section, we assume that monetary policy is nonstochastic and stationary (i.e., $g_t = g = 1$ and $d_t = d$).

Since, in the OLG model, money is the only means of saving, the equilibrium condition is

$$m_t = s_t. \tag{5}$$

Equations (3), (4), and (5) define the equilibrium restrictions of the model. They can be integrated into the equilibrium map

$$\pi_t = \frac{g\int_{[0,1]} \max\{a - b\pi_t^{ke},\ 0\}\mu(dk)}{\int_{[0,1]} \max\{a - b\pi_{t+1}^{ke}, 0\}\mu(dk) - \min\left\{d, \lambda\left[\int_{[0,1]} \max\left\{a - b\pi_{t+1}^{ke}, 0\right\}\mu(dk)\right]\right\}}. \quad \textbf{(6)}$$

We have intentionally left the general notation g to facilitate comparison to subsequent development ($g = 1$ is used in the remainder of this section).

Under perfect foresight, (6) can be written as the difference equation $\pi_{t+1} = R(\pi_t)$. Details are given in Appendix A. Figures 1*A*–*D* show the four possible configurations of the $R(\cdot)$ map (assuming that $1/(1-\lambda) < a/b$). In Figure 1*A* we represent the case in which there are four steady states: π^A, π^B, π^C, and π^O, shown as points A, B, C, and O.[9] Steady states π^A and π^B correspond to the steady states of the standard model and $\pi^C = 1/(1-\lambda)$ to the interior steady state of the constrained regime. Figures 1*B*–*D* show other possible configurations in

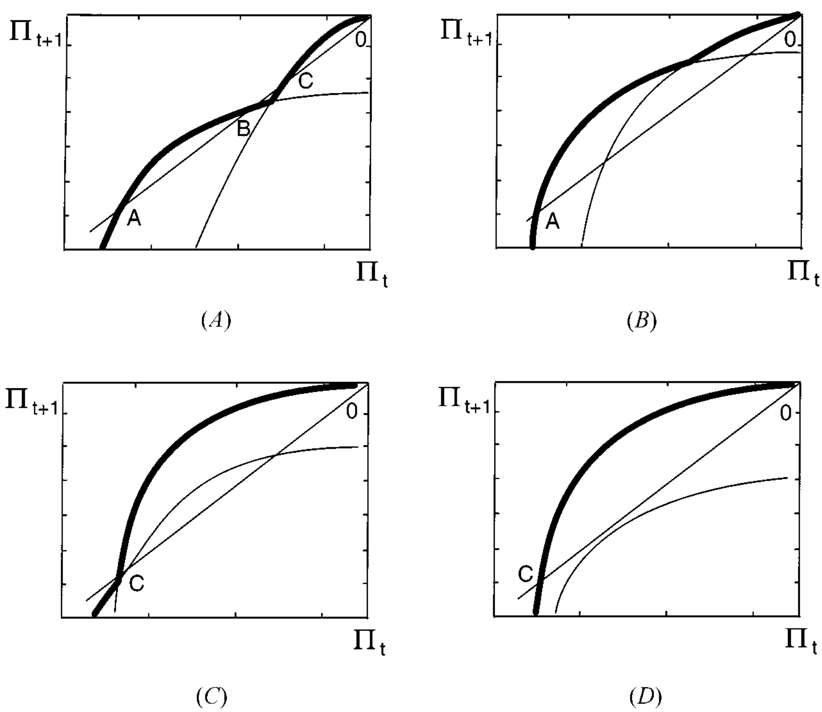

Figure 1. Possible configurations of the $R(\cdot)$ map.

which there are fewer steady states. In Figure 1B the only interior steady state corresponds to the low-inflation steady state of the standard model. In Figures 1C and 1D, there are two steady states corresponding to those of the constrained regime (shown as points C and O). In particular, in Figure 1D, the level of deficit d cannot be sustained as a steady state in the standard model.

In summary, under perfect-foresight dynamics, the addition of a fiscal constraint always introduces a continuum of inflation paths converging to π^O (in some cases, this replaces the paths converging to π^B). Thus, if one focuses on stable perfect-foresight paths, the constraint does not appear to enhance monetary stability. We will see, however, that this conclusion is completely reversed under adaptive learning.

3. EQUILIBRIA WITH LEARNING DYNAMICS

3.1. Homogeneous Case

In place of perfect foresight, we now introduce learning behavior in the formation of expectations and the decisions of individual agents. In forming their expectations of inflation, π^e_{t+1}, the agents are assumed to use an adaptive rule that depends on data, π_{t-1}, and expectations held last period, π^e_t (expectations do not depend on the individual k when agents are homogeneous). At this stage, we introduce an additional assumption to rule out the possibility that aggregate savings are exactly zero even at high inflation rates. Since each young agent has nonnegative savings, aggregate saving can be zero only if the saving of every young agent is exactly zero. Although implied by the model at sufficiently high inflation rates, such a possibility seems implausible and is one that we want to avoid. We therefore assume that actual saving by household k is equal to $s_{t,k} + \epsilon_{k,t+1}$, where $\epsilon_{k,t+1}$ denotes a small nonnegative random variable, iid across k and with positive mean. We remark that, in experimental work on the basic OLG model, Marimon and Sunder (1993, 1994) have found that agents tend to save somewhat more than what an adaptive agent with point expectations and competitive supply would do.

We arrive at the following model in the homogeneous case. After aggregating over identical agents, the savings function takes the form

$$S_t \equiv s_t + \epsilon_{t+1} = \max\{a - b\pi^e_{t+1}, 0\} + \epsilon_{t+1}. \tag{7}$$

Here, $\epsilon_{t+1} = \int \epsilon_{k,t+1}\mu(dk)$ is the aggregate of individual ϵ-savings and, by the law of large numbers, it is nonstochastic and positive. We make the following assumption.

Assumption 1. The sequence ϵ_t is positive and decreasing, with

$$\text{(i)} \quad \lim_{t\to\infty} \epsilon_t = \bar{\epsilon},$$

where $\bar{\epsilon} \geq 0$,

$$\text{(ii)} \quad \lim_{t\to\infty} (\epsilon_{t+1}/\epsilon_t) = 1.$$

The role of the $\epsilon_{k,t}$ is to ensure that currency collapse does not occur in finite time.[10] Various specifications can be made to satisfy Assumption 1. An example is furnished by $\epsilon_t = \epsilon/t$ for any $\epsilon > 0$.

From Section 2, we have the equation determining actual inflation. This takes the form

$$S_t = g_t \frac{S_{t-1}}{\pi_t} + \min[d, \lambda S_t],$$

where $\lambda < 1$; that is,

$$\pi_t = \frac{g_t S_{t-1}}{S_t - \min\{d, \lambda S_t\}}. \tag{8}$$

Here, g_t is allowed to be random, as mentioned in Section 2, with 1 as its mean. Substituting the saving function into this expression, we obtain

$$\pi_t = \frac{g_t \left[\max\left\{a - b\pi_t^e, 0\right\} + \epsilon_t \right]}{\max\left\{a - b\pi_{t+1}^e, 0\right\} + \epsilon_{t+1} - \min\left\{d, \lambda \left[\max\left\{a - b\pi_{t+1}^e, 0\right\} + \epsilon_{t+1} \right] \right\}}. \tag{9}$$

We denote this mapping $\pi_t = F(\pi_{t+1}^e, \pi_t^e, \epsilon_{t+1}, \epsilon_t, g_t)$.

For the learning rule, we start with a simple adaptive rule:

$$\pi_{t+1}^e = \pi_t^e + \gamma_{t+1}\left(\pi_{t-1} - \pi_t^e\right), \tag{10}$$

for $t = 0, 1, 2, \ldots$. For the special case $\gamma_t = 1/t$, π_{t+1}^e is the sample mean of past π_t, which is, of course, a simple form of least-squares learning.

This type of rule is standard in the literature on adaptive learning [see the surveys by Sargent (1993), Marimon (1997), and Evans and Honkapohja (1999)] and the form adapted here is natural for learning steady states in stochastic systems [see, e.g., Evans and Honkapohja (1995)].

We make the following standard assumption.

Assumption 2. The sequence of gains, γ_t, is a positive decreasing sequence satisfying

$$\text{(i)} \quad \sum_{t=1}^{\infty} \gamma_t = \infty,$$

$$\text{(ii)} \quad \sum_{t=1}^{\infty} \gamma_t^p < \infty \ \text{ for some } p,$$

$$\text{(iii)} \quad \lim_{t\to\infty} \sup \left[\frac{1}{\gamma_t} - \frac{1}{\gamma_{t-1}} \right] < \infty.$$

Assumption 2(i) is required to avoid convergence to a non-equilibrium point, and 2(ii) is needed to ensure asymptotic elimination of residual fluctuation in π_t^e. Assumption 2(iii) is needed for technical reasons, but is not usually found to be restrictive. Assumption 2 characterizes a fairly large class of learning rules with decreasing gain. An example is given by $\gamma_t = 1/t$ yielding the case in which π^e_{t+1} is the sample mean.

To obtain the limiting steady states for learning dynamics, we look at

$$G(\pi^e) = \lim_{t\to\infty} F\left(\pi^e, \pi^e, \epsilon_{t+1}, \epsilon_t, 1\right).$$

We show that the possible limiting steady states are then limited to the fixed points of G. It is straightforward to calculate that

$$G(\pi^e) = \begin{cases} \dfrac{(a - b\pi^e)}{a - b\pi^e - d} & \text{if} \quad \pi^e < \dfrac{a - d/\lambda}{b} \quad \text{and} \\ \dfrac{1}{1-\lambda} & \text{otherwise.} \end{cases}$$

There are the following four cases illustrated in Figures 2*A*–*D*. (These figures are drawn to correspond to the different possibilities illustrated in Figures 1*A*–*D*, respectively.)

The standard model with a fixed deficit d corresponds to having $\lambda = 1$. For small d, there are two steady states. In Figure 2*A*, these are shown as points A

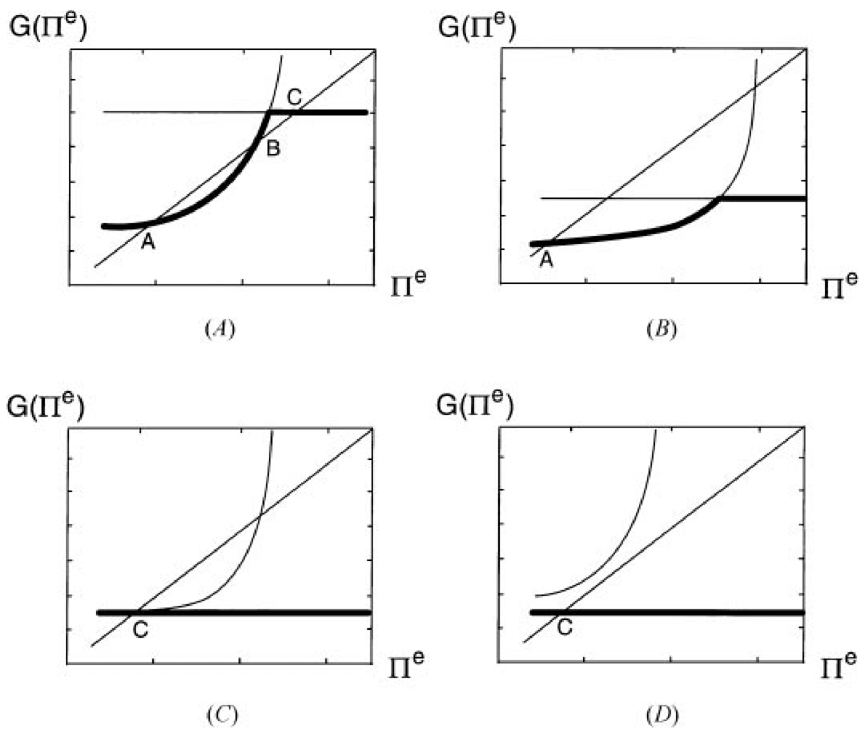

Figure 2. Possible configurations of the $G(\cdot)$ map.

and B and are the intersections of the function $(a - b\pi^e)/(a - b\pi^e - d)$ with the 45-deg line. (Note that $(a - b\pi^e)/(a - b\pi^e - d) \to \infty$ as $\pi^e \to (a - d)/b$.) In our model with $\lambda < 1$, there is a kink to $G(\pi^e)$ at point $(a - d/\lambda)b^{-1}$. After this switch point, the mapping $G(\pi^e)$ becomes horizontal. In Figure 2*A*, the switch occurs above point B, yielding a third equilibrium at point C.

For a lower value of λ, the switch occurs between points A and B. This case results in a unique equilibrium at point A, as shown in Figure 2*B*. Note that in this case the unique solution is the same as the low-inflation steady state of the standard model.

For a low enough value of λ, the switch occurs before the solutions to the standard model. This yields a unique equilibrium on the horizontal part of the mapping $G(\pi^e)$; see Figure 2*C*.

Finally, for high enough values of d, the standard model has no solutions with finite steady-state inflation. In this case, our deficit rule creates a finite steady state for all values $\lambda < 1$. This is illustrated in Figure 2*D*.

The following theorem shows that, under our simple adaptive learning rule, the dynamics converge to the set of locally stable steady states.[11]

THEOREM 1. *Under Assumptions* 1 *and* 2, *as* $t \to \infty$, π_t^e *converges with probability* 1 *to the set*

$$\{\hat{\pi} \mid \hat{\pi} = G(\hat{\pi}),\ G'(\hat{\pi}) < 1\}$$

of stable fixed points of $G(\pi^e)$, *that is, to either point A or C in Figures* 2*A–D.*

$G'(\hat{\pi}) < 1$ is the condition for a steady state $\hat{\pi}$ to be locally stable under the stochastic dynamics of the system (9) and (10). This is a familiar condition in the literature on adaptive learning. For a precise definition of local stability under learning in stochastic frameworks, see Evans and Honkapohja (1998).

The stability properties of π^A, π^B, and π^C under adaptive learning are opposite from their stability under perfect-foresight dynamics.[12] This extends the corresponding results for π^A and π^B in the standard monetary inflation model. Because of the fiscal constraint, the dynamics at high inflation rates drive π_t toward π^C, away from π^O and currency collapse.[13] The fiscal constraint, in the presence of ϵ-saving, prevents even the possibility of currency collapse. (Note that the standard model with learning would remain subject to currency collapse if ϵ-saving were included.) We remark that, provided $1/(1 - \lambda) < a/b$, the paths converging to π^C have positive trading asymptotically, and the deficit can be financed up to the constraint. If $1/(1 - \lambda) > a/b$, then there is still positive trading at all finite times, but asymptotically the economy approaches autarky.

Tightening the fiscal constraint (lowering λ) reduces the maximum possible steady-state inflation rate $\pi^C = 1/(1 - \lambda)$. Moreover, if politically feasible, a sufficient reduction of the value of λ can ensure that the economy converges to the low-inflation steady state π^A. To see this, suppose that the economy is in a high-inflation steady state at point C in Figure 2*A*. If the value of λ can be reduced to a sufficient degree, the steady states B and C in Figure 2*A* will disappear, and the situation becomes qualitatively like Figure 2*B* in which only the steady state with low inflation is present.[14] In this case, the fiscal constraint has helped to stabilize inflation even though it is not binding at the steady state.

3.2. Heterogeneous Learning Behavior

We now allow for heterogeneity in agents' learning behavior. The remainder of our model is unchanged. Thus, aggregate saving, allowing for heterogeneous expectations and the ϵ-shocks, is given by

$$S_t = \int_{[0,1]} \max\{a - b\pi_{t+1}^{ke},\ 0\}\mu(dk) + \epsilon_{t+1}, \tag{11}$$

and realized inflation π_t continues to be given by (8).

We continue to assume that the expectations of agents adapt to the most recent forecast error. In addition, we now allow for the possibilities of inertia in the formation of expectations and random fluctuations in the adaption speeds. In particular, we assume for each agent k the rule

$$\pi_{t+1}^{ke} = \begin{cases} \pi_t^{ke} & \text{with probability } \rho_{t+1} \text{ and} \\ \pi_t^{ke} + \gamma_{k,t+1}\left(\pi_{t-1} - \pi_t^{ke}\right) & \text{with probability } 1 - \rho_{t+1}. \end{cases}$$

The probability $0 \leq \rho_{t+1} < 1$ captures the degree of inertia. The random "gains" $\gamma_{k,t+1}$ are taken to be positive, independent of past information and, for each t, $\gamma_{k,t+1}$ are identically and independently distributed across k. Moreover, $\gamma_{t+1} \equiv E\gamma_{k,t+1}$ is assumed to converge to zero as $t \to \infty$.

The class of learning rules that we consider is, of course, still restrictive. In particular, potentially important cases such as correlated groups of agents and learning by imitation are not permitted. However, our formulation generalizes the standard learning rule (10) in a way that admits heterogeneous learning and still permits us to obtain formal results. As discussed later, the earlier literature on least-squares learning has allowed heterogeneity in the form of different initial priors but has imposed identical learning rules. Theorem 2 is the first in the literature to establish convergence under assumptions that permit agents to process information in different ways.

We replace Assumption 2 with the following assumptions.

Assumption 2′. $\gamma_t(1 - \rho_t)$ is a positive decreasing sequence satisfying

$$\text{(i) } \sum_{t=1}^{\infty} \gamma_t(1 - \rho_t) = +\infty$$

$$\text{(ii) for some } p > 0\text{: } \sum_{t=1}^{\infty} [\gamma_t(1 - \rho_t)]^p < +\infty.$$

$$\text{(iii) } \limsup_{t\to\infty} \left[\frac{1}{\gamma_t(1 - \rho_t)} - \frac{1}{\gamma_{t-1}(1 - \rho_{t-1})}\right] < \infty.$$

Assumption 3. For all k: $\gamma_{kt} \leq \overline{\gamma_t}$, where $\overline{\gamma_t} \to 0$ as $t \to \infty$.

A simple special case that satisfies the above assumptions is averaging with inertia: $\gamma_{kt} = t^{-1}$ and $\rho_{t+1} = \rho$ for some constant $0 \leq \rho < 1$.

We remark that the assumptions on our learning rules imply that even if agents *initially* have homogeneous expectations, heterogeneous adjustments to forecast errors and/or heterogeneous inertia across agents ensure that heterogeneity of expectations will emerge over time. We are nonetheless able to obtain a strong global convergence result for the heterogeneous expectations case, yielding homogeneous expectations in the limit.

THEOREM 2. *Suppose that Assumptions* 1, 2′, *and* 3 *hold and that Condition B*: $g^{\max}/(1-\lambda) < a/b$ *is satisfied. As* $t \to \infty$, π_t^e *converges with probability* 1 *to the set*

$$\{\hat{\pi} \mid \hat{\pi} = G(\hat{\pi}), G'(\hat{\pi}) < 1\}$$

of stable fixed points of $G(\pi^e)$. *Furthermore,* $\forall k$: $\pi_t^{ke} \to \pi_t^e$ *as* $t \to \infty$ *almost surely.*

Condition B in Theorem 2 requires that the upper bound on exogenous money injections, $g^{\max}$, relative to the intensity of the fiscal constraint, $1-\lambda$, be sufficiently low. This condition seems natural, though it is not necessary for avoiding currency collapse in the homogeneous case. However, it is only a small tightening of the condition $1/(1-\lambda) < a/b$, which is needed in either case to ensure that the stable constrained-inflation steady state is bounded away from autarky.

3.3. Relation to the Literature on Learning

Theorem 2 gives a convergence result that is striking in that it gives global convergence almost surely in a model with stochastic shocks and heterogeneous learning rules. We obtain this global convergence result even for cases in which there are multiple steady states.[15]

Heterogeneous expectations formation has been considered only in relatively few papers in the earlier literature. Bray and Savin (1986) consider global convergence of heterogeneous expectations to the rational-expectations equilibrium (REE) by introducing different priors in the context of the cobweb model. However, in their model the equilibrium is unique and the learning rules of different agents are identical. Marcet and Sargent (1989b) analyze some situations of heterogeneity and differential information for a class of models. Finally, a few papers [see, e.g., Frydman (1982) and some papers in the volume edited by Frydman and Phelps (1983)] have explored the issue of dynamics and stability when agents have knowledge of average expectations in the economy.[16]

We go a step further in this direction and introduce heterogeneity of agents' learning rules, within a class of adaptive rules, by incorporating random adaption and inertia as basic behavioral assumptions. We emphasize the way in which these assumptions make the theory more consistent with the experimental data, but potentially they can also play an important role in convergence results. In the context of strategic games, Marimon and McGrattan (1995) show the importance of adaptation and inertia (and experimentation) as behavioral assumptions that guarantee asymptotic convergence results in repeated play.[17]

A different approach to modeling learning is based on genetic algorithms. This approach has been developed in OLG economies by Arifovic (1995, 1998) and Bullard and Duffy (1998). The use of genetic algorithms automatically provides a model of heterogeneous expectations. The genetic algorithm results to date are primarily based on numerical simulations.

In the standard deficits finance model, Marcet and Sargent (1989a), Marimon and Sunder (1993), and Arifovic (1995) found the high-inflation steady state π^B to be unstable under learning, but there do exist some learning rules in which π^B is stable. For nonstochastic models (with homogeneous expectations), Lettau and Van Zandt (1999) show that the usual stability results need not hold if π^e_{t+1} depends on contemporaneous π_t and there is a sufficiently high reaction to it. Because our model allows for stochastic shocks, convergence to an REE requires a decreasing gain $\gamma_t \to 0$; cf. Assumptions 2(ii) or 2′(ii). [See Evans and Honkapohja (1995, p. 200).] This ensures that here π^A is stable and π^B is unstable, in line with the other papers cited above.

As already noted, Bray and Savin (1986) showed global convergence to the unique REE for the cobweb model.[18] For more general models, the methods available in the literature on stochastic approximation [see, e.g., Ljung (1977), Benveniste et al. (1990), and Ljung et al. (1992)] need to be utilized. In the presence of multiple REE, the basic results on global convergence state that (under a compactness condition and under the absence of invariant sets that are more complex than equilibrium points) there is convergence with probability one to the set of equilibria that satisfy a local stability condition. In a model with homogeneous learning rules, Woodford (1990) has used this result when showing the possibility of convergence to sunspot equilibria in the standard overlapping generations model. Here, we use the same technique to derive the set of stable inflationary steady-state solutions to money-financed deficits under heterogeneous learning rules.

Our global convergence results can also be seen as a way to get rid of the "projection facility" condition, used in much of the existing literature that applies stochastic approximation techniques, which has been criticized for its lack of economic foundation [e.g., in Grandmont (1998)].[19] In our model, the "projection facility," is replaced by an economic constraint: Deficits cannot be more than a given fraction of GDP.

4. EXPERIMENTAL EVIDENCE

Previous experimental studies have investigated the standard inflation model. The experiments done by Marimon and Sunder (1993) and Arifovic (1995) show that there is convergence only to the low-inflation and never to the high-inflation steady state. These results accord with the central theoretical predictions of adaptive learning in the standard model. Here, we report on an experiment designed to investigate whether the correspondence extends to the new features that emerge under adaptive learning when the fiscal constraint is introduced.

The theoretical results of the previous two sections can be summarized as follows: First, because of the fiscal constraint, there should be convergence toward some steady-state inflation rate, even for large target deficits. Second, for parameter settings with two stable steady states under learning, one would expect to find path dependence, that is, that the detailed history of the economy would govern the inflation rate to which the economy converges. Third, a tightening of the fiscal constraint reducing the number of steady states ought to shift the economy from a high-inflation to a low-inflation steady state. The goal in this section is to present some experimental evidence in support of these theoretical implications. Obviously, with experiments of finite length, it is impossible to establish the validity of the detailed theoretical assumptions or test for strict asymptotic convergence. However, experimental data can indicate support for the stability properties we have derived and for the implications of policy shifts.

When studying the effect of the various government financial restrictions imposed in historical economies, one must make a priori assumptions about the structure and one must come to terms with the credibility issue. Experimental data are exempt from these problems. First, the underlying structure is known to the experimenter. Second, "the rules of the game" are fixed in advance and never changed ex post (in fact, in contrast with actual governments, the experimenter has little to gain from ex-post changes), leaving little room for credibility problems.

Using the experimental software developed by R. Marimon and S. Sunder, we studied the effect of policy changes in our framework. Appendix B provides a brief description of the experiments.[20] In summary, a network of computers simulates an OLG economy, where experimental subjects participate either as young, old, or outsiders. In the latter case, they are randomly selected to be "born again." Consumption goods cannot be stored between periods and fiat money cannot be stored by old agents. Subjects make predictions of the next period's rate of inflation (knowing prices up to the last period). If they are young, their forecasts are translated by the computer into their optimal savings rates, as in the model in Section 2. These savings, in turn, translate into current consumption and, given the next period's realized inflation, into consumption the following period—when old. These two consumptions, via their endowed preferences, are translated into payoffs (pesetas). If they are not young, then their forecasts are rewarded according to ex-post accuracy. These forecasts of the old and outsiders play no role in the economy, but provide additional information on subjects' forecasting rules. Price is the market-clearing price given current savings of the young, existing money holdings of the old, and the (seignorage) expenditures of the government. Throughout the experiments, lump-sum taxes are zero, and so, the target deficit is identical to target government expenditures. Experimental subjects know the basic parameters of the economy and are informed in advance of any possible policy

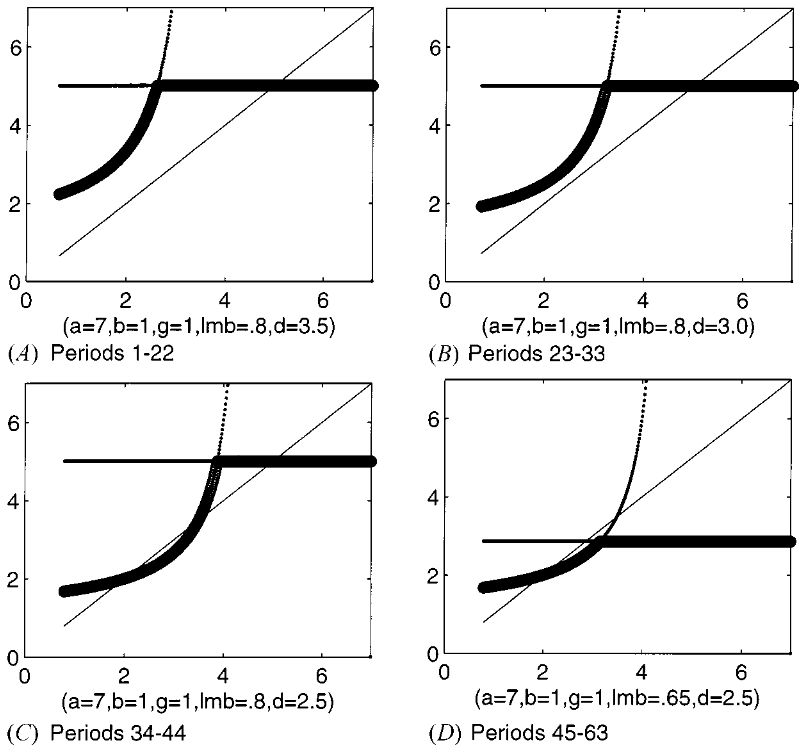

Figure 3. Experimental design of Economy 1.

change. They do not, however, know the exact "termination date." This date is decided by the experimenter and communicated to the subjects after they have submitted that period's forecasts. Therefore, there is a well-defined (endogenous) price for this last period, which is the one that clears the market, given the forecasts. That is, fiat money (of the old) can be translated into consumption units, even though commodity markets do not open in the last period.

Two experimental economies were conducted in two experimental sessions at the Universitat Pompeu Fabra experimental lab (LeeX, July 15, 1994). Figure 3 describes Economy 1 (as in Figure 2, the thick line represents the G map). This economy has four regimes, corresponding to three changes of parameters. The new parameters were announced two periods before the changes took place. Starting from a situation of unsustainable desired government expenditures, the per-capita deficit d is reduced (in period 23), but since the new level is still unsustainable, the structure of steady-state equilibria remains unchanged.[21] A second change (in period 34), to a sustainable level of government expenditures, expands the set of stationary equilibria from one to three. Finally, the restriction on government expenditures λ is made more stringent (in period 45), reducing the set of stationary equilibria from three to one. Our aim, with such a design, was to see whether we

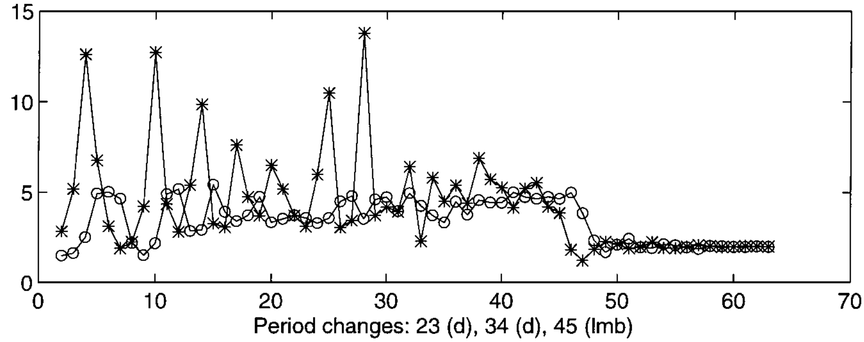

(*A*) Inflation Economy 1:* = Realized, O = mean prediction

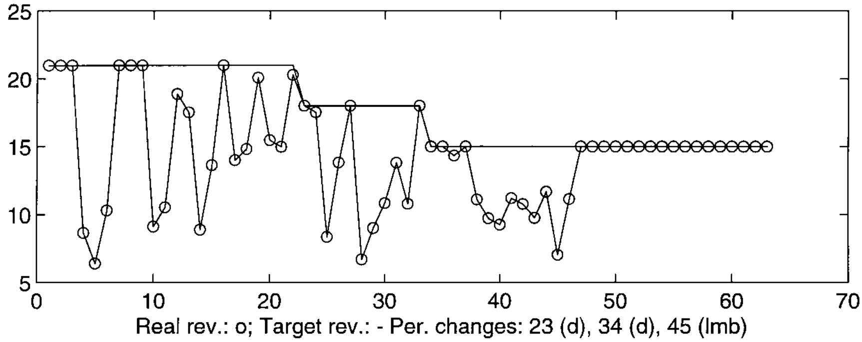

(*B*) Government real revenues in Economy 1

Figure 4. Results for Economy 1 experiment.

would observe convergence first to the high-inflation steady state, and then, after the change of government's restriction, to the low-inflation steady state.

Figures 4 and 5 show the experimental results. As can be seen, the deficit is unsustainable (up to period 33), but inflation is very high and very volatile.[22] The first reduction of the deficit, although it does not change the high steady-state (gross) inflation rate (which is 5 from period 1 to period 44), may have some influence on the inflation path. After period 33, the deficit is within the constraint for some periods and, even when it is binding, inflation fluctuates around the stable high-inflation steady state. The volatility of inflation is somewhat reduced in this regime. Although in this regime there are two stable steady states, the history of the economy has led the agents to the high-inflation equilibrium.

The strengthening of the fiscal constraint (in period 45) completely modifies the inflation path dynamics. Although (gross) inflation is 4.1 in period 44, and remains at 3.8 in period 45, it falls to 1.8 in period 46, remaining around the low-inflation steady state after that (see Figures 4*A* and 5*D*).

The direct effect of the fiscal constraint λ can be seen in Figure 4*B*.[23] All periods in which the revenues are below the targeted revenues are periods in which the fiscal constraint has been binding. As can be seen, the constraint is never binding

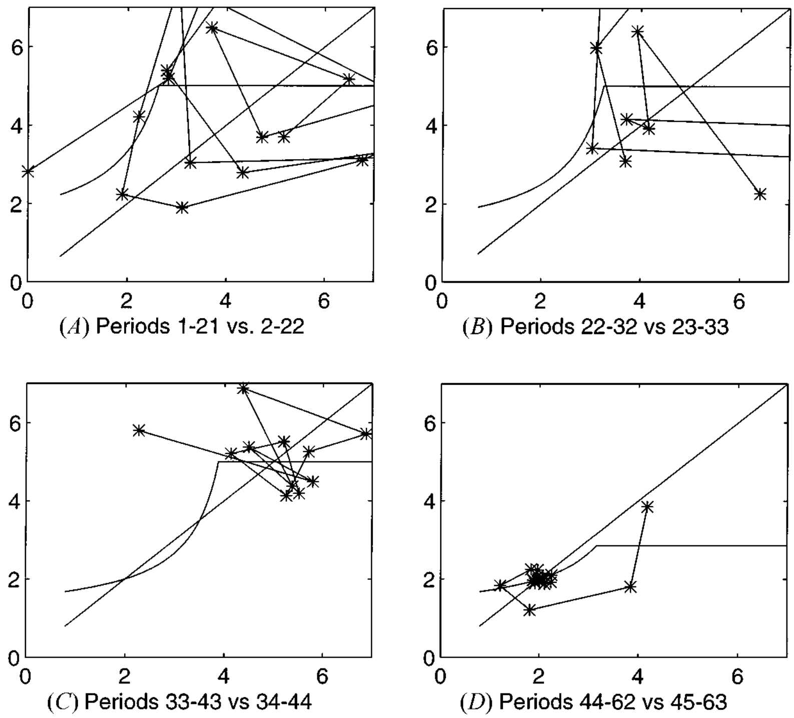

FIGURE 5. Results for Inflation in Economy 1 by period.

after period 47.[24] Thus, in this experiment, a tightening of the fiscal constraint led to a situation in which the constraint is not binding in the end. This apparently paradoxical result is a consequence of the reduction in the number of steady states after the policy change. Under adaptive learning, the economy is driven to the low-inflation steady state at which the target revenues are attained.

In Economy 2, we studied the change from a policy regime in which currency collapse is almost possible [i.e., condition $1/(1-\lambda) < a/b$ is not satisfied] to a regime in which the tighter fiscal constraint ensures positive saving in the constrained steady state. Figure 6 shows the experimental design of Economy 2. For the first nine periods, the fiscal constraint, λ, is 0.9 and it is reduced to 0.8 in period 10. That is, after period 10, there are three stationary equilibria with positive saving.

The results of Economy 2 (Figure 7) are very interesting and show how the stabilization mechanism can be very strong. It is worthwhile following the sequence of events in some detail. In the first six periods, inflation is very high and volatile. Forecasted inflation is high,[25] resulting in low savings. In particular, in period 6, the average predicted inflation, for period 7, by the six members of the young generation is 6.91,[26] which results in an effective expenditure of 0.04 per capita (as opposed to the planned 0.5), which is low relative to the expenditure of the previous

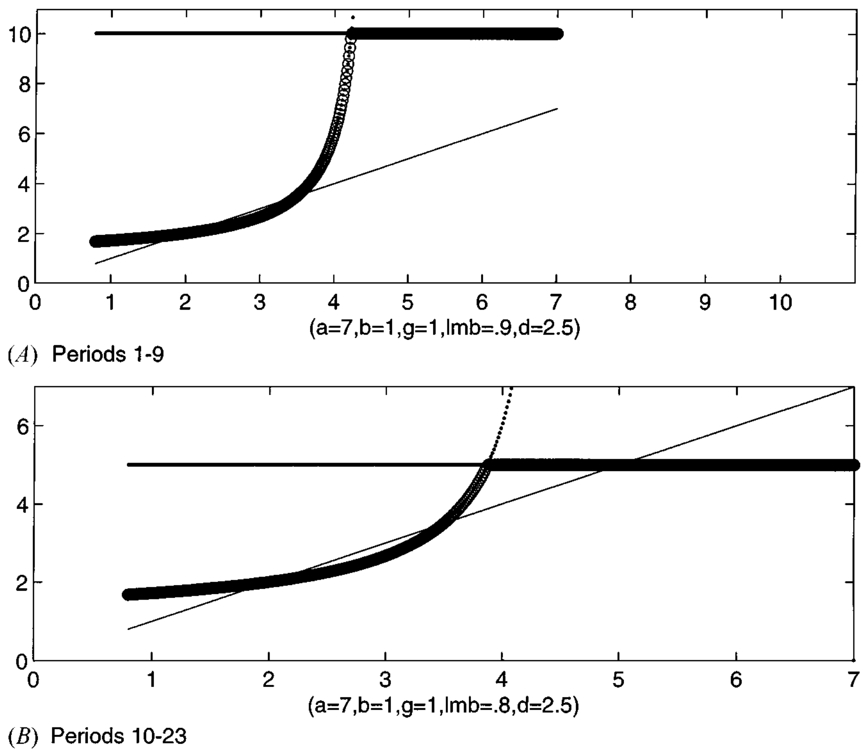

Figure 6. Experimental design of Economy 2.

period, and results in a very high inflation, $\pi_6 = 109.5$. [Recall that, by equation (8), $\pi_t = 10(S_{t-1}/S_t)$, as long as the constraint, $\lambda = 0.9$, is binding.] However, the somewhat lower price expectations of next period's young (6.12, formed before the jump of p_6 was observed) result in a dramatic reduction of inflation next period ($\pi_7 = 1.02$).

Nothing prevents inflation from jumping high again the following period. However, in period 8, the announcement is made that there will be a strengthening of the fiscal constraint in period 10. The subjects have had previous experience (Economy 1) of the reduction of the inflation rate that followed such a policy change. The predicted mean inflation by the young in period 8, for period 9, is 3.53, substantially less than in previous periods. As a result, the fiscal constraint is not binding in period 8 (see Figure 7*B*) and inflation remains low. After that, it seems the inflation path appears to be characterized by the stability properties of the low-inflation stationary equilibrium.

These data also show how "adaptive" subjects are able to "anticipate" the effect of an announced change of policy regime once they have had experience of a similar change.[27] In Economy 2, this results in convergence to the low-inflation stationary state, though the increased fiscal discipline did not imply such a result. Inflation

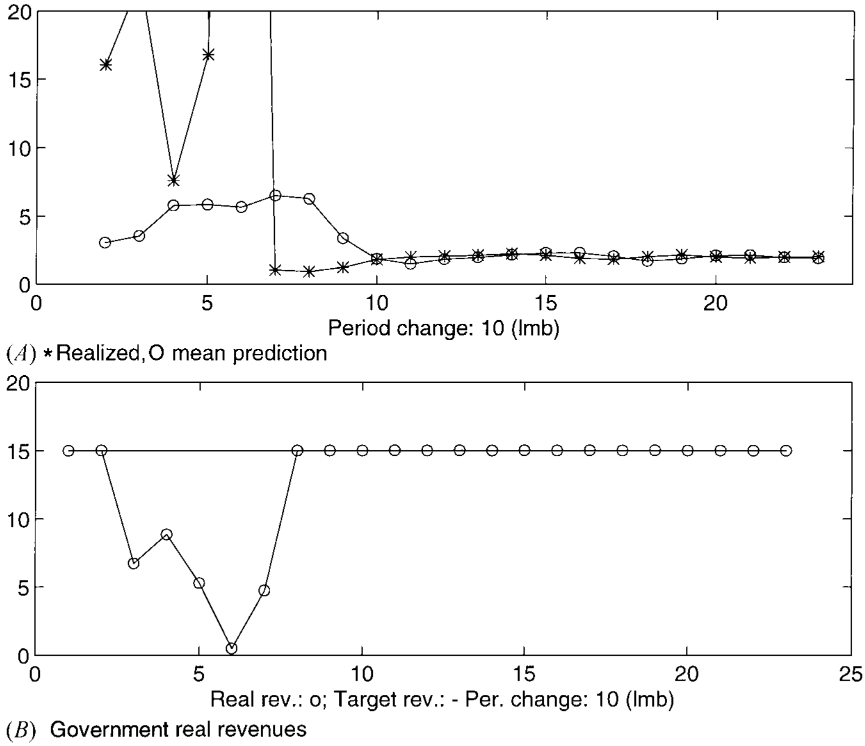

FIGURE 7. Results for Inflation in Economy 2.

rates could have converged to—or stayed around—the high-inflation steady state as in Figure 5*C*. Note that the policy parameters are identical for Figure 3*C* (experiment 1) and Figure 6*B* (experiment 2). Although the theory in Sections 3 and 4 does not allow for anticipation effects, the different outcomes clearly reflect the path dependence of the model under learning. In principle, the theory could be formally extended to allow agents to condition their forecasts on such announcements.

5. CONCLUDING REMARKS

The economic message from this paper is twofold. First, a limitation on the size of the public deficit as a fraction of a measure of aggregate market activity can be a mechanism for the avoidance of explosive price paths, leading to currency collapse, which might take place in the absence of such a constraint, most notably when real target deficits are large. Second, a sufficient tightening of the fiscal constraint can ensure convergence to the low-inflation steady state despite the existence of multiple steady states in the absence of the fiscal constraint. Note that in the latter situation a discrete change in the fiscal constraint generally is required: The response to the parameter change is nonlinear, and disappearance of some steady states is needed. Learning dynamics then do a significant part of the work in the adjustment to the new inflation equilibrium. The striking contrast between the rational expectations and learning analyses of the role of the fiscal constraint points to the importance of the study of learning processes in policy design.

This paper also makes a significant step toward studying economies with adaptive heterogeneous agents: We obtain formal global convergence results in a stochastic nonlinear model while allowing for the heterogeneity of expectations that is apparent from experiments. By introducing random adaption speeds and inertia, we have allowed for much more heterogeneity than has the previous theoretical literature. Further generalizations would be possible. A finite number of different consumer types could be accommodated by appropriately stacking the variables and the dynamic system describing the learning dynamics.

Finally, we have presented some supporting experimental evidence. Although our (limited) experimental results cannot test our detailed formal assumptions, they do support the idea that a fiscal constraint on monetary policy can have a major stabilizing role. Our experimental data are consistent with the view that agents behave adaptively in forming their expectations and that policies have the effects predicted by our model with adaptive learning.

NOTES

1. In related work, Marcet and Nicolini (1998) study the use of external constraints, taking the form of convertibility rules, for inflation stabilization. They develop their results in the context of Latin American inflationary episodes. Monetary instability under learning of pure interest rate control in a cash-in-advance model has been shown by Howitt (1992).

2. The discussion of this paragraph somewhat oversimplifies the results on adaptive learning in the standard model. The literature on learning is discussed in Section 3.3.

3. The fiscal constraint is binding only in the transition to the low-inflation steady state. It does not affect the steady state.

4. The monetary asset in OLG models serves solely as a store of value and can be interpreted as a nominal one-period bond. Setting the coupon rate equal to zero for convenience yields its interpretation as money. See Farmer (1993, Ch. 6). Furthermore, in OLG models, money and real bonds are perfect substitutes, as discussed for example by Sargent (1987, Ch. 7, Sect. 4); the crucial variable determining the equilibrium price path is the total liability of the government. By extending the model along the lines of Sargent and Wallace (1981), it would be possible to allow for the simultaneous existence of money and nominal bonds at the cost of added theoretical complexity.

5. Each agent k of the initial generation lives only in period 1 and is endowed with $\omega = \omega^2$ of the consumption good. He also has an endowment of fiat money of h_0 and his preferences are represented by $u_k(c_0) = \ln(c_0)$.

6. As noted later, in equilibrium we also have s_t equal to the real money stock. Recall that, in a steady state, real balances are the inflation tax base, with the tax rate given by the net inflation rate. More general setups than our simple OLG model would allow us to distinguish between the inflation tax base, savings, and GDP.

7. Government expenditures can also be financed in part by lump-sum taxes, provided they are constant over time. For convenience, these taxes are set at zero.

8. Since in our framework deficits are equal to seignorage, the constraint can also be interpreted as a limitation on seignorage. If the government can issue debt, combined fiscal constraints on debt and deficits imply constraints on seignorage. Note that both types of constraints (i.e., constraints on deficits and direct constraints on seignorage) are in place in some countries, for example, in the EMU.

9. Steady-state $\pi^O = a/b$ corresponds to the autarky solution in which money is not valued. There are also nonstationary perfect-foresight paths converging to either π^B (point B) or π^O (point O).

10. This could also be achieved by introducing heterogeneous endowments or tastes. Provided some agents have a strong saving motive, aggregate saving will remain positive for a wide range of expected inflation rates. A limit $\bar{\epsilon} > 0$ is formally equivalent to having a class of agents with no second-period endowment.

11. The proofs of the theorems are in Appendix A.

12. Recall also that, under perfect foresight, the dynamics at high inflation rates follow paths converning toward π^O. This would remain true under perfect foresight if we introduced ϵ-savings (convergence would be to a small perturbation of π^O if $\bar{\epsilon} > 0$).

13. In a different framework, Evans and Ramey (1995) also show the possibility of stable high rates of inflation if agents' adaptation is costly. However, they show that currency collapse will result for other parameter values in their model.

14. Confining the analysis of the model to rational-expectations dynamics might mislead one to view such a tightening as harmful since, after the reduction in λ, the situation is qualitatively like Figure 1*B* with dynamic paths approaching the autarky point O.

15. Recently, Bischi and Marimon (1999) characterized the global stability properties of a (deterministic) model similar to the one studied here, allowing constant gain learning rules [i.e., $\gamma_t = \gamma \in (0, 1)$]. Their approach, based on global bifurcation theory, differs from ours and reinforces the conclusions of our work.

16. Townsend (1983) considers similar issues for rational learning.

17. The "eductive" approach to learning suggested by Guesnerie (1992) also emphasizes heterogeneous expectations.

18. Evans and Honkapohja (1998) provide a generalization of the Bray and Savin result to simultaneous equation models with a unique REE.

19. The "projection facility" amounts to constraining expectations to a compact set in the basin of attraction of a locally stable steady state. For another way of avoiding the projection facility, see Evans and Honkapohja (1998).

20. See, for example, Marimon and Sunder (1993, 1995) for a more detailed account of the experimental approach followed here.

21. An unsustainable target deficit is one that precludes the existence of a steady state in the model without the additional government financial constraint.

22. The volatility in this regime is not surprising in light of equation (8) or (A.4). With $\lambda = 0.8$ and $g = 1$, $\pi_t = 5(S_{t-1}/S_t)$ if the fiscal constraint binds. Futhermore, if inflation expectations are sufficiently high, $S_t = \epsilon_t$ so that inflation is heavily influenced by the random ϵ-saving.

23. In Figures 4*B* and 7*B*, the vertical axis shows total government revenues, i.e., $n \times d$, where in these experiments $n = 6$.

24. The average revenues by subperiods are as follows: Periods 1–22, 15.9; Periods 23–33, 13.2; Periods 34–44, 12.1; and Periods 45–63, 14.4. This shows that revenues are still high with the final (stabilizing) policy, even though it is less expansionary than the ones previously used (e.g., in Periods 34–44).

25. As is explained in Appendix B, reported forecasts have been truncated to be less than $a/b = 7$.

26. In Figure 7*A*, the mean prediction corresponds to the mean prediction of all the subjects—not only the young—which in period 7 is 6.50.

27. Similar "anticipation effects" are reported by Marimon and Sunder (1994) and in their more recent ongoing work.

28. If $1/(1 - \lambda) > a/b$, this branch intersects the 45-deg line from below at the autarky point a/b.

29. See Marcet and Sargent (1989c, Appendix) or Woodford (1990, Appendix) for statements of some versions of Ljung's theorems.

30. In Economies 1 and 2: $N = 14$ and $n = 6$.

31. This procedure of imposing the reward structure effectively induces the same responses as if the agents had log utility; see the discussion of induced value theory in Friedman and Sunder (1994, pp. 12–15).

32. Translated from the Spanish instructions which had a somewhat less technical language (but the same equations, etc.). As usual, the instructions were explained in detail to the subjects.

33. This was not a good feature of the software used in the experiment because it constrained final observed forecasts. The current version follows the above formulas, allowing for arbitrarily large forecasts (up to computer limitations!). Nevertheless, this feature of the forecasting procedure did not have any impact on the experimental economy.

REFERENCES

Arifovic, J. (1995) Genetic algorithms and inflationary economies. *Journal of Monetary Economics* 36, 219–243.

Arifovic, J. (1998) Stability of equilibria under genetic algorithm adaption: An analysis. *Macroeconomic Dynamics* 2, 1–21.

Benveniste, A., M. Metivier & P. Priouret (1990) *Adaptive Algorithms and Stochastic Approximations*. Berlin: Springer-Verlag.

Bischi, G.I. & R. Marimon (1999) Global Price Stability with Adaptive Agents and Fiscal Constraints. Mimeo, European University Institute.

Bray, M.M. & N.E. Savin (1986) Rational expectations equilibria, learning, and model specification. *Econometrica* 54, 1129–1160.

Bruno, M. (1989) Econometrics and the design of economic reform. *Econometrica* 57, 275–306.

Bullard, J. & J. Duffy (1998) On learning and the stability of cycles. *Macroeconomic Dynamics* 2, 22–48.

Evans, G.W. & S. Honkapohja (1995) Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models. *Econometrica* 63, 195–206.

Evans, G.W. & S. Honkapohja (1998) Economic dynamics with learning: New stability results. *Review of Economic Studies* 65, 23–44.

Evans, G.W. & S. Honkapohja (1999) Learning dynamics. In J.B. Taylor & M. Woodford (eds.), *Handbook of Macroeconomics*, Vol. 1, pp. 449–542. Amsterdam: Elsevier.

Evans, G.W. & G. Ramey (1995) Expectation calculation, hyperinflation and currency collapse. In H. Dixon & N. Rankin (eds.), *The New Macroeconomics: Imperfect Markets and Policy Effectiveness*, pp. 307–336. Cambridge, UK: Cambridge University Press.

Farmer, R. (1993) *The Economics of Self-fulfilling Prophecies*. Cambridge, MA: MIT Press.

Fischer, S. (1984) The economics of Israel. In K. Brunner & A.H. Meltzer (eds.), *Monetary and Fiscal Policy and Their Applications*, supplement to the *Journal of Monetary Economics* 20, 7–52.

Friedman, D. & S. Sunder (1994) *Experimental Methods, A Primer for Economists*. Cambridge, UK: Cambridge University Press.

Frydman, R. (1982) Towards an understanding of market processes: Individual expectations, learning, and convergence to rational expectations equilibrium. *American Economic Review* 72, 652–668.

Frydman, R. & E.S. Phelps (eds.) (1983) *Individual Forecasting and Aggregate Outcomes, "Rational Expectations" Reexamined*. Cambridge, UK: Cambridge University Press.

Grandmont, J.-M. (1998) Expectations formation and stability of large socioeconomic systems. *Econometrica* 66, 741–781.

Guesnerie, R. (1992) An exploration of the eductive justifications of the rational-expectations hypothesis. *American Economic Review* 82, 1254–1278.

Howitt, P. (1992) Interest rate control and nonconvergence to rational expectations. *Journal of Political Economy* 100, 776–800.

Judd, K. (1985) The law of large numbers with a continuum of IID random variables. *Journal of Economic Theory* 35, 19–25.

Lettau, M. & T. Van Zandt (1999) Robustness of Adaptive Expectations as an Equilibrium Selection Device. Mimeo, INSEAD.

Ljung, L. (1977) Analysis of recursive stochastic algorithms. *IEEE Transactions on Automatic Control* AC-22, 551–575.

Ljung, L., G. Pflug & H. Walk (1992) *Stochastic Approximation and Optimization of Random Systems*. Basel: Birkhauser.

Marcet, A. & J.P. Nicolini (1998) Recurrent Hyperinflations and Learning. Working paper 1875, CEPR.

Marcet, A. & T.J. Sargent (1989a) Least squares learning and the dynamics of hyperinflation. In W.A. Barnett, J. Geweke, & K. Shell (eds.), *Economic Complexity: Chaos, Sunspots, Bubbles, and Nonlinearity*. Cambridge, UK: Cambridge University Press.

Marcet, A. & T.J. Sargent (1989b) Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy* 97, 1306–1322.

Marcet, A. & T.J. Sargent (1989c) Convergence of least-squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, 337–368.

Marimon, R. (1997) Learning from learning in economics. In D.M. Kreps & K.F. Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications,* Vol. I, pp. 278–315. Cambridge, UK: Cambridge University Press.

Marimon, R. & E. McGrattan (1995) On adaptive learning in strategic games. In A. Kirman & M. Salmon (eds.), *Learning and Rationality in Economics*, pp. 61–101. Oxford: Basil Blackwell.

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

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

Marimon, R. & S. Sunder (1995) Does a constant money growth rule help stabilize inflation? Experimental evidence. *Carnegie-Rochester Conference Series on Public Policy* 43, 111–156.

Sargent, T.J. (1987) *Dynamic Macroeconomic Theory*. Cambridge, MA: Harvard University Press.

Sargent, T.J. (1993) *Bounded Rationality in Macroeconomics*. Oxford: Oxford University Press.

Sargent, T.J. & N. Wallace (1981) Some unpleasant monetarist arithmetic. *Federal Reserve Bank of Minneapolis Quarterly Review* 5(3), 1–17.

Sargent, T.J. & N. Wallace (1987) Inflation and the government budget constraint. In A. Razin & E. Sadka (eds.), *Economic Policy in Theory and Practice*. London: Macmillan.

Townsend, R.M. (1983) Forecasting the forecasts of others. *Journal of Political Economy* 91, 546–588.

Woodford, M. (1990) Learning to believe in sunspots, *Econometrica* 58, 277–307.

APPENDIX A: DERIVATIONS AND PROOFS

A.1. PERFECT-FORESIGHT CASE

Equation (6) describes the equilibrium dynamics of the economy, actual inflation as a function of agents' expected inflation for the current and the following period. When expectations are homogeneous, this equation is of the form

$$\Phi\left(\pi^e_{t+1}, \pi^e_t, \pi_t\right) = 0.$$

We can close the equilibrium condition by postulating the *rational-expectations hypothesis*. That is, $\pi^{ke}_t = \pi_t$. Then, rational-expectations equilibrium paths are given by a difference equation,

$$\pi_{t+1} = R(\pi_t),$$

which takes the precise form

$$\pi_{t+1} = \begin{cases} \left(\dfrac{a-d}{b}+1\right) - \dfrac{a}{b}\dfrac{1}{\pi_t} & \pi_{t+1} \le \dfrac{a-d/\lambda}{b} \\ \left(\dfrac{a}{b}+\dfrac{1}{1-\lambda}\right) - \dfrac{a}{b}\dfrac{1}{(1-\lambda)\pi_t} & \pi_{t+1} > \dfrac{a-d/\lambda}{b} \end{cases}. \tag{A.1}$$

Notice that, as long as $\pi_{t+1} \le (a-d/\lambda)b^{-1}$, the $R(\cdot)$ map corresponds to the standard equilibrium map of models of seignorage, where, if the level of the deficit, d, is sustainable, there are two steady states (π^A, π^B); the one with higher inflation, π^B, is usually identified with a steady state of hyperinflation. On the other hand, another concave map is defined for the range in which $\pi_{t+1} > (a-d/\lambda)b^{-1}$. Provided $1/(1-\lambda) < a/b$, this map has two steady states, the low $\pi^C = 1/(1-\lambda)$, and the high $\pi^O = a/b$. At this last steady state, savings are zero and money ceases to have value.[28] When they exist, all four steady states define stationary rational-expectations equilibria. Furthermore, there is a continuum of nonstationary rational-expectations equilibrium paths: In the long run, these paths are either at π^B or in the nonmonetary equilibrium corresponding to π^O.

A.2. LEARNING WITH HOMOGENEOUS EXPECTATIONS

Here we prove Theorem 1.

Proof of Theorem 1. We set up the model as a recursive stochastic algorithm and use the results of Ljung (1977). Substituting for π_{t-1} in the adaptive learning rule, we get

$$\pi_t^e = \pi_{t-1}^e + \gamma_t \left[F\left(\pi_{t-1}^e, \pi_{t-2}^e, \epsilon_{t-1}, \epsilon_{t-2}, g_{t-2}\right) - \pi_{t-1}^e\right].$$

To apply the results of Ljung (1977), we need to set up the system in two parts. First, we define an equation for the evolution of a vector of state variables φ_t. We thus let

$$\varphi_t = A\left(\pi_{t-1}^e\right)\varphi_{t-1} + B\left(\pi_{t-1}^e\right)e_t,$$

where $\varphi_t = (\varphi_{1,t}, \varphi_{2,t}, \varphi_{3,t})'$, $e_t = (1, g_{t-2})'$, and

$$A\left(\pi_{t-1}^e\right) = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \qquad B\left(\pi_{t-1}^e\right) = \begin{bmatrix} 0 & 0 \\ \pi_{t-1}^e & 0 \\ 0 & 1 \end{bmatrix}.$$

The second part of the algorithm is the equation for the vector of parameters (here simply the scalar π_t^e) taking the form $\pi_t^e = \pi_{t-1}^e + \gamma_t Q(t, \pi_{t-1}^e, \varphi_t)$, where

$$Q\left(t, \pi_{t-1}^e, \varphi_t\right) \equiv \left[F\left(\pi_{t-1}^e, \varphi_{1,t}, \epsilon_{t-1}, \epsilon_{t-2}, \varphi_{3,t}\right) - \pi_{t-1}^e\right].$$

It is straightforward to verify that Ljung's assumptions B are met.[29] Moreover, $\pi_t^e \le N_\delta \equiv \max\{a/b, g_{\max}/(1-\lambda)\} + \delta$, for any $\delta > 0$, infinitely often. We prove this by contradiction. If $\pi_t^e > N_\delta$ for $\forall t > T_0$ for some T_0, then

$$\pi_t \le \frac{g_t \epsilon_t}{(1-\lambda)\epsilon_{t+1}} \quad \forall t > T_0.$$

Thus, from equation (10), it follows that π_t^e is a strictly decreasing sequence and, since $\sum_{t=1}^{\infty} \gamma_t = \infty$, it must eventually fall below N_δ. It follows that $\pi_t^e \leq N_\delta$ infinitely often.

The assumptions of Theorems 1 and 2 of Ljung (1977) are now satisfied. It follows that, with probability 1, π_t^e converges to the stable fixed points of the associated differential equation

$$\frac{d\pi^e}{d\tau} = G(\pi^e) - \pi^e. \qquad \blacksquare$$

A.3. LEARNING WITH HETEROGENEOUS EXPECTATIONS

We can rewrite the learning rule

$$\pi_{t+1}^{ke} = \pi_t^{ke} + \xi_{t+1}^k \gamma_{k,t+1}\left(\pi_{t-1} - \pi_t^{ke}\right), \tag{A.2}$$

where for each agent k the variable ξ_{t+1}^k is a Bernoulli random variable, independent of past information, and independent across k, which takes the value 0 with probability ρ_{t+1} and 1 with probability $1 - \rho_{t+1}$. In addition, $\forall j, k : \xi_{t+1}^j$ are assumed independent of current π_t^{ke} and $\gamma_{k,t+1}$.

Integrating equation (A.2) over k, we get

$$\pi_{t+1}^e = \pi_t^e + \int \xi_{t+1}^k \gamma_{k,t+1}\left(\pi_{t-1} - \pi_t^{ke}\right)\mu(dk).$$

We have

$$\int \xi_{t+1}^k \gamma_{k,t+1} \pi_t^{ke} \mu(dk) = \left\{ \int \xi_{t+1}^k \mu[\mu(dk)] \right\} \left[\int \gamma_{k,t+1} \pi_t^{ke} \mu(dk) \right] + \underset{\mu(dk)}{\mathrm{Cov}}\left(\xi_{t+1}^k, \gamma_{k,t+1}\pi_t^{ke}\right).$$

Using the law of large numbers for continua of random variables [see Judd (1985)], $\int \xi_{t+1}^k \mu(dk) = 1 - \rho_{t+1}$ and, from the independence assumption,

$$\underset{\mu(dk)}{\mathrm{Cov}}\left(\xi_{t+1}^k, \gamma_{k,t+1}\pi_t^{ke}\right) = 0.$$

Moreover, we get

$$\int \gamma_{k,t+1} \pi_t^{ke} \mu(dk) = \left[\int \gamma_{k,t+1} \mu(dk) \right] \left[\int \pi_t^{ke} \mu(dk) \right] = \gamma_{t+1} \pi_t^e.$$

Analogously, using independence and the law of large numbers, we have

$$\int \xi_{t+1}^k \gamma_{k,t+1} \pi_{t-1} \mu(dk) = (1 - \rho_{t+1}) \gamma_{t+1} \pi_{t-1}.$$

These arguments yield

$$\pi_t^e = \pi_{t-1}^e + \gamma_t (1 - \rho_t)\left(\pi_{t-2} - \pi_{t-1}^e\right). \tag{A.3}$$

We have arrived at an aggregate learning rule for expectations of inflation that is very similar to equation (10) even though we have allowed for heterogeneous expectations.

However, with heterogeneous expectations, aggregate savings, which continues to determine inflation through equation (8), is no longer given by equation (7) but instead by (11). ■

Proof of Theorem 2. We introduce the notation $M_t(\pi) = \{k \mid \pi_t^{ke} < \pi\}$, $M_t = M_t(a/b)$ and $\mu_t(\pi) = \mu[M_t(\pi)]$, where μ denotes the Lebesgue measure. Then, inflation is given by

$$\pi_t = \frac{g_t \cdot \left[\int_{M_t} \left(a - b\pi_t^{ke}\right)\mu(dk) + \epsilon_t\right]}{\int_{M_{t+1}} \left(a - b\pi_{t+1}^{ke}\right)\mu(dk) + \epsilon_{t+1} - \min\left\{d, \lambda \cdot \left[\int_{M_{t+1}} \left(a - b\pi_{t+1}^{ke}\right)\mu(dk) + \epsilon_{t+1}\right]\right\}}. \quad \textbf{(A.4)}$$

We first show that $\exists T$ and $\delta > 0{:}\, \forall t > T{:}\, \pi_t \leq a/b - \delta$.

From Condition B, $\exists \delta > 0{:}\, \pi_t > a/b - \delta$ implies $\pi_t > [g^{\max}/(1-\lambda)](1+\kappa)$ for some $\kappa > 0$ and any t. On the other hand, (A.4) implies

$$\pi_t \leq \left(\frac{g^{\max}}{1-\lambda}\right)\left(\frac{s_{t-1} + \epsilon_t}{s_t + \epsilon_{t+1}}\right).$$

Therefore,

$$\frac{s_{t-1} + \epsilon_t}{s_t + \epsilon_{t+1}} > 1 + \kappa.$$

Using Assumption 1(ii), it follows that

$$\exists T_1 > 0, \qquad 0 < \kappa_1 < \kappa{:}\ \forall t > T_1{:}\ \frac{s_{t-1}}{s_t} > 1 + \kappa_1. \quad \textbf{(A.5)}$$

On the other hand, for any t,

$$\frac{s_{t-1}}{s_t} = \frac{\int_{M_t(\pi_{t-1})} s_{t-1,k}\mu(dk) + \int_{M_t(\pi_{t-1})^c} s_{t-1,k}\mu(dk)}{\int_{M_t(\pi_{t-1})} s_{t,k}\mu(dk) + \int_{M_t(\pi_{t-1})^c} s_{t,k}\mu(dk)},$$

where $M_t(\pi_{t-1})^c = [0,1]\backslash M_t(\pi_{t-1})$, the set-theoretic complement. Next, note that $k \in M_t(\pi_{t-1})^c$ implies $s_{t,k} \geq s_{t-1,k}$, since π_{t+1}^{ke} is a convex combination of π_t^{ke} and π_{t-1}. Moreover, using the learning rules (A.2), we have

$$s_{t,k} = \max\left[s_{t-1,k} - b\xi_{t+1}^k \gamma_{k,t+1}\left(\pi_{t-1} - \pi_t^{ke}\right), 0\right],$$

so that, for all t large enough, $k \in M_t(\pi_{t-1})$ implies $s_{t,k} \geq s_{t-1,k} - b\bar{\gamma}_t \pi_{t-1}$, where $\bar{\gamma}_t$ is given in Assumption 3. It follows from Assumptions 1(ii) and 3 that

$$\frac{s_{t-1}}{s_t} \leq \frac{a\mu_t(\pi_{t-1}) + \int_{M_t(\pi_{t-1})^c} s_{t-1,k}\mu(dk)}{\int_{M_t(\pi_{t-1})^c} s_{t-1,k}\mu(dk) - \mu_t(\pi_{t-1})b\bar{\gamma}_t\pi_{t-1} + \int_{M_t(\pi_{t-1})^c} s_{t-1,k}\mu(dk)} \to 1, \quad \text{as } t \to \infty. \quad \textbf{(A.6)}$$

Statements (A.5) and (A.6) are contradictory, so that $\pi_t < a/b - \delta$ for some $\delta > 0$ small enough and all t sufficiently large.

Using (A.2), even the slowest π_t^{ke} will reach a value $\leq a/b$ in finite time. Thus, $\forall t > \bar{T}$ sufficiently large, $\pi_t^{ke} \leq a/b$. Since

$$\int \pi_t^{ke} \mu(dk) = \pi_t^e,$$

it follows from (A.4) that

$$\pi_t = \frac{g_t\left(a - b\pi_t^e + \epsilon_t\right)}{a - b\pi_{t+1}^e + \epsilon_{t+1} - \min\left[d, \lambda\left(a - b\pi_{t+1}^e + \epsilon_{t+1}\right)\right]} = F\left(\pi_{t+1}^e, \pi_t^e, \epsilon_{t+1}, \epsilon_t, g_t\right) \tag{A.7}$$

from $\bar{T}$ onward. Combining equations (A.3) and (A.7), we arrive at

$$\pi_t^e = \pi_{t-1}^e + \gamma_t(1 - \rho_t)\left[F\left(\pi_{t-1}^e, \pi_{t-2}^e, \epsilon_{t-1}, \epsilon_{t-2}, g_{t-2}\right) - \pi_{t-1}^e\right] \tag{A.8}$$

to which we can again apply the method of proof in Theorem 1. The argument is analogous to that in Section 3.

Finally, consider the individual expectations. For each k, we consider the bivariate recursive algorithm in the variables (π_t^e, π_t^{ke}). To obtain the equation for π_t^{ke}, we substitute (A.7) into equation (A.2), yielding

$$\pi_t^{ke} = \pi_{t-1}^{ke} + \gamma_t(1 - \rho_t)\left(\frac{\xi_t^k}{1 - \rho_t}\right)\left(\frac{\gamma_{kt}}{\gamma_t}\right)\left[F\left(\pi_{t-1}^e, \pi_{t-2}^e, \epsilon_{t-1}, \epsilon_{t-2}, g_{t-2}\right) - \pi_{t-1}^{ke}\right]. \tag{A.9}$$

The system consisting of equations (A.8) and (A.9) is now in standard form and satisfies Ljung's assumptions. The associated differential equation for (π_t^e, π_t^{ke}), when linearized, has the coefficient matrix

$$\begin{pmatrix} G'(\hat{\pi}) - 1 & 0 \\ G'(\hat{\pi}) & -1 \end{pmatrix}$$

at an equilibrium $(\hat{\pi}, \hat{\pi})$. Since this matrix is stable, the result follows. ■

APPENDIX B: DESCRIPTION OF EXPERIMENTS

These experiments were conducted on student subjects at the Universitat Pompeu Fabra in Barcelona. $N \geq 2n$ subjects were recruited for each session.[30] Of these, n subjects each played the role of "young" and "old" generations, respectively, in any given period, while the remaining $N - 2n$ subjects waited as interested onlookers. At the beginning of each period, n subjects were randomly picked from the group of those who were not young in the preceding period, and the remaining $N - 2n$ were left in the waiting pool. This process ensured that every subject had to wait a random number of periods (minimum 0) between exiting the economy and reentering it as young again.

After reading and explaining the instructions (instructions for Economy 1 follow this narrative), the subjects participated in five to six periods of a trial economy. Fiat money was labeled "francs" and the consumption good was labeled "chips." The number of chips "consumed" were converted into Spanish pesetas at the end of the exit period of each subject. Total pesetas accumulated in this manner were paid to subjects at the end of the session in cash. Most sessions lasted for about three hours, and each subject took home 2,000–3,000 peseta on average from each session.

All subjects were asked to predict inflation in the price of chips, $\tilde{\pi}_{t+1} = \tilde{p}_{t+1}/\tilde{p}_t$, at the beginning of each period t.

From the individual inflation forecasts, the computer constructed an optimal supply for each individual, and the market supply from the individual supplies.[31] The computer also calculated a market demand function for chips from the money balances of the old after taking into account the government policies on fiscal deficit and money growth. The computer calculated the market-clearing price and allocations, and distributed this information to all subjects. When the experimenter terminated an economy without advance warning, franc balances of the young in the last period were converted into chips at the market-clearing price calculated for the following (nonplayed) period, before the announcement of termination was made.

B.1. INSTRUCTIONS FOR ECONOMY 1 (JUNE 1994[32])

This is an experiment in decision making. The Ministry of Education of Spain has provided funds for this research. The instructions are simple; follow them carefully. The money you earn depends on the decisions you and others make. You will make decisions with the help of the computer. This money will be paid to you in cash at the end of the experiment.

This experiment is divided into many periods. Your role may change from period to period. You will have the opportunity to buy "chips," sell "chips," and make predictions of what will happen in the future. The attached Information and Record Sheet will help you keep a record of your decisions and determine their value to you.

The type of currency used in this market is francs. The only use of this currency is to buy and sell chips. It has no other use. The money you take home with you is in pesetas. The procedures for determining the number of pesetas you take home with you is explained later in these instructions.

You will participate in a market for two consecutive periods at a time. Let us call the first of these periods your *entry* period (because you begin your participation then), and the second your *exit* period (because you end your participation in the market). Different individuals may have different entry and exit periods. We shall tell you when you *enter* and *exit* the market. You may enter and exit more than once depending on the number of periods for which the market is operated.

B.2. Trading and Recording Rules

(1) All entry-period players are sellers and all exit-period players (and possibly the experimenter) are buyers. At the beginning of the entry or exit period you will receive an amount of chips (endowment). This endowment always will be greater in your

entry period (young) than your exit period (old). You cannot carry the chips from one period to the next.

(2) Every exit-period player (old) pays all his francs to entry-period players (young) in exchange for chips at a market price determined in the manner explained below.

(3) At the beginning of each period, *every* player (young, old, and outsider) must state the prediction of price ratio for the following period: $\tilde{p}_{t+1}/\tilde{p}_t = 1 +$ inflation rate. Predictions of the entry players (young) will be used to determine the number of chips they wish to sell according to the formula given later on.

(4) After considering the francs available from the exit players (old), offers made by entry players (young) and experimenter's policy (government) about financing the debt with francs and/or incrementing the quantity of francs in circulation, the market-clearing price is computed and announced. Exit players (old) and the experimenter pay this price for each chip they buy. Each entry player (young) will be informed of the number of chips he/she has been able to sell at the market price, and each exit-period player (old) will be told of the number of chips that he/she has been able to buy with his/her francs on hand.

(5) Each exit (old) and outside player receives a reward one period later, depending on how close his/her prediction of price ratio is to the actually observed price ratio. At the end of that period, the experimenter announces the most accurate predicted price ratio, and the market price ratio.

(6) After transaction information is received, through the computer, each entry player (young) can compute the chips remaining on hand (consume). The francs received from sale will be used to buy chips in the exit period that follows immediately. You can carry your francs on hand forward to the exit period by entering them in the column "Francs" in the Record Sheet.

(7) Each exit player (old) records the number of chips purchased on the Information and Record Sheet. Then the experimenter computes the pesetas earned by using formula (B.1) or (B.1′) given below. This amount is the profit of the exit-period player (old) who records this profit on the Information and Record Sheet. At the end, the experimenter will pay each player the total amount of profits in pesetas.

(8) The experimenter may terminate the market at any time. Without any announcement in advance, the participants will be informed which is the last period of the experiment; francs held by all entry-period players (young) are converted into chips using the market-clearing price of the following period.

(9) At the end of the experiment, add up the earnings and prediction rewards columns of your Information and Record Sheet. The experimenter will pay you the sum of these (your cumulative earnings) in pesetas.

B.3. Payoffs

The number of "pesetas" you will earn to take home with you for any pair of entry–exit period will be

$$\max\{0, e(\log c_1 + \log c_2)\}, \qquad \textbf{(B.1)}$$

where e, the exchange rate of "utility to pesetas," is set to 11 pesetas/chip today.

This means that the greater the product of chips you consume at the end of your entry and exit periods, the greater the number of pesetas you earn to take home with you.

The first period of the market will be an entry period for some of you (as described above). For some of you, however, this first period itself will be an exit period and you will receive the exit period endowment (ω^2). In addition, each of you for whom the first period is an exit period will receive an amount of francs from the experimenter at the beginning of this period. These participants have to use all these francs to buy chips during the exit period because the francs you hold at the end of an exit period are worthless; they cannot be converted into pesetas directly. The payoff of such individuals at the end of Period 1 will be

$$e(\log c_2) \tag{B.1$'$}$$

At the beginning of each period, all of you will be asked to predict the price ratio of chips for the next period. For example at the beginning of period 1, you will be asked what will be the ratio of the price of chips in period 2 to the price of chips in period 1 ($\tilde{p}_2/\tilde{p}_1$). Please note that, at the time you are asked to predict their ratio, you do not know either of the two prices and, for example, if the price ratio (now) is 1.10, this means that the inflation rate is 10%.

B.4. Automatic "Chip" Supply

If you are in your entry period, the computer will use the price ratio you enter ($\tilde{p}_{t+1}/\tilde{p}_t$) to compute the number of chips you sell in that period, as follows:

$$\max_{0 \le s \le \omega^1} \left\{ 0, \log(\omega^1 - s) + \log\left(\omega^2 + \frac{s}{\tilde{p}_{t+1}/\tilde{p}_t} \right) \right\} + \epsilon$$

This means that your optimal decision when you are convinced the price ratio will be $\tilde{p}_{t+1}/\tilde{p}_t$ is

$$s_{i,t} = \max\left\{0, 0.5\left[\omega^1 - (\tilde{p}_{t+1}/\tilde{p}_t)\omega^2\right]\right\} + \epsilon \tag{B.2}$$

Notice that the higher the price ratio you enter the fewer the number of chips you will sell. If your price ratio prediction is ω^1/ω^2 , you will sell ϵ chips, where $\epsilon \ge 0$ is a very small random number representing small "precautionary savings." If your price ratio prediction is higher than ω^1/ω^2, the computer will ask you to submit a new prediction.[33]

B.5. Prediction Rewards

The price ratio predicted by the exit and outside players will earn them a reward, the more accurate your prediction, the greater the reward (determined by their absolute magnitude of relative error in your prediction):

$$\max\left\{0, e_2\left(1 - \left|\left[\frac{(p_{t+1}/p_t) - (\tilde{p}_{t+1}/\tilde{p}_t)}{(p_{t+1}/p_t)}\right]\right|\right)\right\}, \tag{B.3}$$

where ($\tilde{p}_{t+1}/\tilde{p}_t$) is your prediction. Today, $e_2 = 20$.

This means that the exact prediction will earn you 20 pesetas, and prediction error will reduce your reward proportionately until it becomes zero (0) for an error of 100% or more. Your prediction reward cannot be negative. Entry players will not participate in these prediction rewards.

B.6. Policy Parameters and Changes

At the beginning of an economy, you will be told the amount of government expenditures planned to be financed by printing money ("Govt. deficit" in your screen), as well as, the restriction imposed on the government ("Govt. restriction" in your screen), that is the maximum fraction of total savings that can be used for government consumption. For example, if "Govt. restriction" is 0.9 and total savings are $S_t = 12$, then total government expenditure (deficit) cannot be more than $0.9 \times 12 = 10.8$, even if the announced "Govt. deficit" was 11.

In some economies, these "Govt." parameters may be changed. Nevertheless, any change will either be announced at the outset or, at least two periods in advance (with a message in your screen).