Exp Econ (2010) 13: 146–166
DOI 10.1007/s10683-010-9233-9

An experimental test of Taylor-type rules with inexperienced central bankers

Jim Engle-Warnick · Nurlan Turdaliev

Received: 20 November 2007 / Accepted: 4 January 2010 / Published online: 11 February 2010

Abstract We experimentally test monetary policy decision making in a population of inexperienced central bankers. In our experiments, subjects repeatedly set the short-term interest rate for a computer economy with inflation as their target. A large majority of subjects learn to successfully control inflation by correctly putting higher weight on inflation than on the output gap. In fact, the behavior of these subjects meets a stability criterion. The subjects smooth the interest rate as the theoretical literature suggests they should in order to enhance stability of the uncertain system they face. Our study is the first to use Taylor-type rules as a framework to identify inflation weighting, stability, and interest-rate smoothing as behavioral outcomes when subjects try to achieve an inflation target.

Keywords Monetary policy · Taylor rule · Experimental economics · Repeated games

JEL Classification C91 · E42

Electronic supplementary material The online version of this article (http://dx.doi.org/10.1007/s10683-010-9233-9) contains supplementary material, which is available to authorized users.

J. Engle-Warnick (✉)
Department of Economics, McGill University, 855 Sherbrooke St. W., Montreal QC H3A 2T7, Canada
e-mail: jim.engle-warnick@mcgill.ca

N. Turdaliev
Department of Economics, 401 Sunset Avenue, University of Windsor, Windsor ON N9B 3P4, Canada
e-mail: nurlan@uwindsor.ca

1 Introduction

The importance of monetary policy decision rules, which are frameworks within which central banks are committed to make their decisions, has grown rapidly over the last decade. A purpose of adopting such rules is to provide central bankers with more effective and robust monetary policy (through better and more stable actions) and to help the public better understand central banks' actions (through better transparency of those actions). The growing importance of monetary policy rules is apparent when one considers the amount of research conducted in this area (e.g., Taylor 1999a; Bernanke and Woodford 2005), and the influence of this research on policy makers.

The Taylor rule (Taylor 1993), which linearly maps the inflation gap (the difference between inflation and a given inflation target) and output gap (the difference between output and potential output) into the central bank's instrument (say, the short-term nominal interest rate), is an example of a simple policy rule. Substantial research has been done in investigating whether the Taylor rule is an accurate description of the behavior of central banks (see, e.g., Clarida et al. 1999), and the Taylor rule has inspired very active theoretical research of the properties of simple policy rules, such as optimality and robustness (Svensson 1997b; Levin et al. 1999).

From a historical perspective, the Taylor rule provides a useful framework to examine monetary policy over time. Using this approach, Taylor (1999b) assesses several periods in history of U.S. monetary policy, identifies turning points in monetary policy, and quantifies the size of past mistakes as well as the degree of effectiveness of different policy rules. John Taylor provides the following summary for this research (Taylor 1999a): simple policy rules behave nearly optimally and are more robust than complex rules.

The simplest Taylor-type rule is

$$i_t = \gamma_0 + \gamma_1(\pi_t - \bar{\pi}) + \gamma_2 x_t,$$

where i_t is the central bank's instrument chosen at time t, π_t is the inflation rate, $\bar{\pi}$ is a given inflation target, x_t is the output gap (the difference between log-output and log-potential output). The coefficients γ_1 and γ_2, i.e., the weights placed on inflation and output, are positive. Taylor suggested $i_t = 4 + 1.5(\pi_t - 2) + 0.5\, x_t$ in his original paper.

In this paper, we present a new use for Taylor-type rules: as a lens through which to analyze the monetary policy decisions of inexperienced central bankers. In our economics experiments, inexperienced 'central bankers' are asked to conduct monetary policy in a laboratory to maximize 'social welfare', which is directly related to their payoff. Subjects repeatedly set the short-term interest rate in two simple economies in which Taylor-type rules are optimal.

We find that Taylor-type rules explain a significant share of the variance of the data, particularly when we fit them on a subject-by-subject basis, thus allowing for heterogeneity with respect to the individual weights subjects place on inflation and output. Indeed, most of our subjects learn how to control the economy: we find that their behavior meets the stability criterion for the resulting dynamical system.[1] By contrast, the estimate for the unsuccessful subjects indicates that they failed to learn how to meet the stability criterion.

The experimental laboratory is the ideal environment in which to study central bank decision making. Because we know the model and its parameters, as well as the inflation target (Judd and Rudebusch 1998), we are able to derive a stability criterion and to determine whether the estimated coefficients imply stability of the economic system: we find that they do. We are also able to determine the importance of applying sufficient weight on inflation by determining if subjects' behavior is optimal, conditional on the weight they are applying to inflation: we find that it is. Because of the limited and controlled amount of data available to the subjects, not subject to later revisions (Orphanides 2001), or measurement error as with the output gap, we can more confidently identify whether or not subjects gradually change, or smooth, the interest rate: we find that they do.

Our laboratory finding of smoothing corresponds to empirical evidence that central banks smooth as well. Clarida et al. (1999) have sparked interest in the issue by identifying it as an important phenomenon. Sack (1998) studies a central bank facing uncertainty with regard to the model of the economy: interest-rate smoothing arises because of the gradual learning of the slope of the IS curve (Rudebusch 2001 also presents an uncertainty model). Since the intercept of the IS curve is stochastic, it is optimal for the central banker to smooth, given her loss function.[2] Our experiment has this flavor because subjects are not told the specification of the economic model, but are given the opportunity to learn by repeatedly playing the game. While the optimal Taylor-type rules in the two economies do not involve a lagged instrument, i.e. they do not exhibit interest-rate smoothing, the smoothing behavior exhibited by the subjects is likely related to their uncertainty about the economy they face.[3]

In practice, it seems that such model uncertainty is an important concern for central bankers. Referring to the Fed's lack of knowledge as to how the economy works, the then Vice-Chairman of the Board of Governors of the Federal Reserve System Alan Blinder (1998) said:

[1] For an example of such a criterion, consider a particularly simple version of our environment when coefficients $\alpha_1 = 1, \beta_1 = 0$, and $\hat{\alpha}_1 = 0$ in the system (3)–(4). As mentioned in Taylor (1999c), it is important to have the inflation response coefficient γ_1 larger than 1 ('stability threshold'). Taylor attributes a better monetary policy in the 1980s and 90s to γ_1 being above 1 as opposed to the 1960s when it was below 1.

[2] Other reasons for smoothing include Cukierman (1996), who argues that such a policy makes the financial system more stable; Caplin and Leahy (1997), who advocate the view that a central bank may avoid too frequent and large changes in the interest rate because of its concern of being judged as poorly informed or making bad decisions; Goodfriend (1991), who puts forward an argument that interest rate smoothing is a better framework for a central bank to communicate its policy to financial markets; Woodford (2003), who demonstrates that an interest-rate smoothing policy may have a beneficial effect in an environment with a forward-looking private sector by anchoring its expectations regarding future policy. We are interested in studying this issue in the simplest possible environment, and leave testing these other explanations for future research.

[3] Informally, a cautious monetary policy of gradual (and, in a way, predictable) changes in interest rate is referred to as *interest-rate smoothing*. More formally, a monetary policy exhibits interest-rate smoothing if the corresponding Taylor rule has a significant coefficient on the lagged interest rate.

> "What can you do to try to guard against failure? ... First of all, be cautious. Don't oversteer the ship. If you yank the steering wheel really hard, a year later you may find yourself on the rocks."

There have been at least two other related studies of central bank decision making in the laboratory. Arifovic and Sargent (2003) find that policy makers can find ways to achieve a time-inconsistent optimal inflation rate in an expectational Phillips curve model. In their experiment policy makers set the inflation target and consumers report inflation expectations. Blinder and Morgan (2005) study the advantages of committee decision making in a problem more similar to ours: their question is whether groups or individuals identify and respond to a shock to money demand better. Our study may shed additional light on these existing results by specifically studying the decision rules that subjects use while determining monetary policy. Not only do our central bankers smooth the interest rate, they also put proper weights on inflation and output to be within the stability region and successfully manage the economy.

2 Experimental design and procedures

2.1 The macroeconomic model

We study a dynamic stochastic general equilibrium model, a variant of the standard New Keynesian model (see, e.g., Clarida et al. 1999). The following two aggregate equations describe the economy:

$$x_t = q\,x_{t-1} + (1-q)E_t x_{t+1} - b(i_t - E_t\pi_{t+1}) + u_t, \tag{1}$$

$$\pi_t = r\,\pi_{t-1} + (1-r)a_1 E_t\pi_{t+1} + a_2 x_t + v_t. \tag{2}$$

Here x_t, π_t, i_t denote correspondingly output gap, inflation, and interest rate in period t. The two other terms u_t and v_t are shocks. The first equation indicates a negative relationship between the output gap x_t and the expected real interest rate $i_t - E_t\pi_{t+1}$ (it is sometimes referred to as an "IS" curve). It is obtained by log-linearizing the household Euler equation. The second equation is a Phillips curve that shows a positive relationship between inflation and the output gap. It can be derived from an environment of monopolistically competitive firms. The equation is obtained by log-linearizing around the steady state of the aggregate firms' pricing decision.

Lagged output gap and inflation are added on the basis of empirical evidence (Fuhrer and Moore 1995) and theoretical work on adjustment costs and adaptive expectations. In the particular case of $q = 1$ and $r = 1$ we obtain the backward-looking model studied in Svensson (1997a). This type of model does not take into account agents' responses to changes in policies, and thus assumes that the model parameters do not change as policy changes.[4] The advantage is the great simplification of the

[4] We can usefully think of our model as analogous to a long line of models in experiments in Industrial Organization, in which human subjects set prices in a game against computer robot buyers. In these experiments, the goal is to learn as much as possible about seller behavior by isolating it apart from the buyers. In our experiment, we wish to learn as much as possible about the policy maker, and we plan to add human consumers in future work.

analysis obtained by treating the private sector's expectations as an adaptive process. After some arrangements we can obtain the following model:[5]

$$x_{t+1} = \beta_1 x_t - \beta_2 (i_t - \pi_t) + \nu_{t+1}, \tag{3}$$

$$\pi_{t+1} - \bar{\pi} = \alpha_1 (\pi_t - \bar{\pi}) + \hat{\alpha}_1 (\pi_{t-1} - \bar{\pi}) + \alpha_2 x_t + \epsilon_{t+1}. \tag{4}$$

In general, the coefficients α_2 and β_2 are assumed to be positive, and all others are nonnegative; in addition, $\beta_1 < 1$. In this economy, inflation is serially correlated and increasing in lagged output. Output is serially correlated and decreasing in the lagged short-term real interest rate $i_t - \pi_t$. In this dynamic economy $\bar{\pi}$ is the long-run inflation rate and the long-run output gap is zero (when optimal monetary policy is pursued).[6]

The monetary authority's objective is to choose a sequence of interest rates $\{i_\tau\}_{\tau=t}^{\infty}$ to minimize the following intertemporal loss function:

$$E_t \sum_{\tau=t}^{\infty} \delta^{\tau-1} \frac{1}{2} (\pi_\tau - \bar{\pi})^2, \tag{5}$$

where $\bar{\pi}$ is the inflation target. We assume that δ is close to 1, i.e. the central bank is sufficiently patient.

The central bank's policy rule can be written as

$$i_t = \gamma_0 + A(L)\pi_t + B(L)x_t,$$

where $A(L)$ and $B(L)$ are lag polynomials (note that we suppress $\bar{\pi}$ hereafter). The lower the degree of the polynomials A and B the 'simpler' is the policy rule. The optimal policy rules in our models have the following form:

$$i_t = \gamma_0 + \gamma_1 \pi_t + \hat{\gamma}_1 \pi_{t-1} + \gamma_2 x_t,$$

where $\hat{\gamma}_1 = 0$ if $\hat{\alpha}_1 = 0$.

2.2 The experimental design

We ran two versions of the model in the experimental laboratory. In Model 1 only the first lag of inflation affects current inflation ($\hat{\alpha}_1$ is zero), and in Model 2 we distribute the effect of inflation over an additional lag. For Model 1 we used Weymark's (2004) coefficient estimates of the Canadian economy, rounding down β_1 to ease the demands of tricky dynamics on our subjects. We chose these coefficients because we wanted the system dynamics to be somewhat consistent with a real economy, and

[5] For more details about this kind of models, see Clarida et al. (1999), Fuhrer and Moore (1995), and Ball (1997).

[6] Our model generalizes Svensson's model, in which inflation is assumed to have a unit root. In Svensson's model inflation and output are increasing in an exogenous variable, which provides a potentially interesting addition for further study.

Table 1 Coefficients for Models 1 and 2 based on Canadian economy

Coefficient	Description	Model 1	Model 2	Weymark
α_1	First lag inflation on inflation	0.50	0.20	0.4964
$\hat{\alpha}_1$	Second lag inflation on inflation	0.00	0.30	–
α_2	First lag output on inflation	0.15	0.15	0.1324
β_1	First lag output on output	0.90	0.90	0.9386
β_2	First lag instrument on output	0.75	0.75	0.7311

we wanted to leave the door open for future study of open economy issues. The coefficients are presented in Table 1. We set both the potential level of output and the long-run inflation rate at 5.00, we set the inflation target at 5.00 and we always started the system in the steady state.

It can be demonstrated that the optimal solution for Model 1 is a Taylor-type rule, and for Model 2 a linear rule that adds non-zero weight to the first lag of inflation. As is shown in the on-line appendix, the optimal rule for Model 1 is:

$$i_t = 5.00 + 3.22\pi_t + 1.87x_t, \tag{6}$$

and the optimal rule for Model 2 is:

$$i_t = 5.00 + 5.00\pi_t + 0.533\pi_{t-1} + 1.60x_t. \tag{7}$$

While we did not expect subjects to apply optimal weights to the relevant variables, our design does make the following testable predictions: (1) from Model 1 to Model 2, the relative weight on inflation should increase, and (2) in Model 2 but not in Model 1, non-zero weight should be placed on the first lag of inflation.

For inexperienced subjects, this is a very difficult decision-making problem. Imagine you are balancing a broom with the handle in the palm of your hand and the broom head up in the air. The location of the head of the broom is analogous to inflation and the location of your hand the instrument; your target is a vertical broom. Now imagine the head of the broom starts to move away from you (i.e., inflation increases past the target). You must move your hand past the physical location of the broom head to bring it back towards you (i.e., you must set the instrument higher than the rate of inflation), and as the broom moves towards you, you must bring your hand back with it (i.e., you must reduce the instrument as inflation nears the target). Moving the palm of the hand past the head of the broom is analogous to placing a weight on inflation greater than one. It is not obvious that the feedback subjects receive while playing the game will be as useful as the visual feedback one receives balancing the broom.

2.3 Experimental procedures

Since our subjects were college students,[7] it was necessary for us to deal with heterogeneity of their knowledge of macroeconomics in some way. We decided not to tell them that they would control a laboratory economy, but we gave them opportunities to practice and learn how the economy worked.[8] Thus our design presents a tough test because it does not allow subjects to bring prior information about monetary policy to bear on their decision making: subjects were not told the decisions they were making had anything to do with the economy, and they were not shown the equations driving the system. Instead they were presented with "chip levels in two containers labelled Container A and Container B", told that the levels were related to each other and the instrument, that increasing the instrument would tend to lower the chip levels, and that each period the relationships were computed and randomness was added to each chip level. Container A actually corresponded to output and Container B to inflation. The goal was to keep the chip level in Container B (i.e, inflation) as close to 5.00 as possible. Subjects were told that the relationships would not change during the entire session, and that the randomness was independent and identical in each period.[9]

Upon arriving at the experimental laboratory, subjects were seated in front of a computer screen. Their decision-making consisted of entering a number for the instrument with up to two decimal places, and clicking an OK button. The screen revealed the entire history of the game, including all past realizations of inflation, output, and all past instrument values. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

Subjects were able to practice the game for no pay before playing it once for pay (as in the schooling choice problem in Houser et al. 2004). The practice sessions were set up so that there were five 10-period games, followed by two 25-period games, followed by 50-period games. Subjects were told that they could practice as often as they liked, with the limitation that the lab was booked for an hour and a half. If subjects wished to play the game for pay after practicing only 10-period games, we suggested they try at least a 25-period game first to improve their chance of earning money; performance in 10-period games is not a good predictor for 50-period games because it may take longer than 10 periods to lose control of the system. When they did play the game for pay, subjects were told that all the rules and relationships remained

[7]The subject pool is drawn from English speaking students in Montreal. The candidate universities in Montreal are McGill University, Concordia University, University of Montreal, and University of Quebec at Montreal.

[8]This design choice causes our results to not be comparable with Blinder and Morgan (2005) who revealed to the subjects that the committee problem was to conduct monetary policy. Our design follows Houser et al. (2004), who kept the context of their schooling choice problem hidden from the subjects. We chose this design to reveal heterogeneity with regard to the ability to control an economic system, rather than to reveal heterogeneity regarding prior knowledge of how a macro economy works. This design allows the economic incentives of the game to guide learning how to play the game, and causes all subjects to begin the game with as close to the same knowledge of how to play the game as possible. Our conjecture is that putting the game in context would speed up learning for subjects familiar with macroeconomics, and increase the average weight placed on inflation.

[9]These informational conditions regarding the specific economic model were similar to those in Blinder and Morgan (2005).

the same, except that the game for pay would last either 51 or 52 periods, and that they would not be told which.[10] To the subjects in every session, either outcome was equally probable; this breaks up an end-game strategy, and we only analyze data up to fifty periods. All subjects in both treatments experienced exactly the same shocks in the game for pay, which were drawn from the normal distribution with mean zero and standard deviation 0.15 using the random number generator in Ox (Doornik 2002).

We paid subjects using a scheme similar to that used in Arifovic and Sargent (2003). Each period, a subject's loss was converted to "period points" (say, P_t) by adding 0.10 to the negative of the computed loss.[11] For subjects who were doing reasonably well, this transformed period points would be a positive number. The subjects earnings (say, E) were \$25.00 times the ratio of the sum of all their period points to the sum of the maximum possible period points they could have earned:

$$E = 25.00 \sum_{t=1}^{50} P_t / 5.0,$$

where $P_t = 0.1 - 0.5(\pi_t - \bar{\pi})^2$. It can be shown that this payoff scheme provides the same incentive as minimizing losses in (5) provided δ is close to 1. (Intuitively, when $\delta = 1$, the intertemporal payoff for the first 50 periods is the sum of one-period payoffs without discounting: $\sum_{t=1}^{50} P_t$. The expression for E above is exactly this save that it is multiplied by 5.) Thus the maximum theoretical earnings for decisions were \$25.00, if no losses ever occurred. We have discovered through practice that it is possible to earn \$23.00. Although earnings may be displayed as negative, subjects never earned less than the show-up fee. The display presented the period loss, period points, and the earnings that would be made if the current period were the last period of play.

We ran the experiments at the Bell University Laboratory in Electronic Commerce and Experimental Economy in Montreal in the winter term of 2004, where English speaking subjects were recruited from four local universities. The subjects earned a \$10.00 show-up fee (which covers the cost of public transportation to travel to the lab, which is not located on campus, for many of our subjects). Sixty-eight subjects participated with Model 1 earning an average of \$26.68, and seventy subjects participated with Model 2 earning an average of \$26.60, including the show-up fee. Subjects varied by the amount of practice games they played, and were paid and dismissed as they finished. Each subject played at her/his own pace. Sessions did not last beyond an hour and a half. We analyze only the games for pay.[12]

[10] For programming purposes we flipped a coin before the first session and used that result to run 52-period games.

[11] Thus, we convert the problem of minimization of loss into the problem of maximization of payoff.

[12] Because of the difficulty of the decision making problem, we ran two pilot sessions, which we used only to ensure that subjects could bring the economy under control. We determined this simply by observing the number of subjects who made nonnegative earnings in the pilots; we did not analyze the data, and the data are available upon request. We learned that two variables are important with regard to complexity of the problem: the coefficient on inflation, which cannot be a unit root because it makes the dynamics too difficult, and the variance of the shocks, which hide the underlying structure of the economy from the decision maker.

3 Results

Figure 1 shows the time series of inflation, output gap, and instrument choice if the central bank uses the optimal rule in Model 1 in our experiments. This simulation is valid for every subject because every participant experienced the same sequence of shocks in the games for pay. The figure shows that it is possible to limit inflation (the solid line) between 4.5% and 5.5% by selecting instrument values roughly in the range from 4 to 6.

Figure 2 shows the same series for Model 2. The more complicated dynamics require a slightly wider range of instrument selection, but again one can see that inflation can be kept within a very narrow band of control. Due to the way we distributed the lagged effect of inflation and due to the identical sequence of shocks, one can also

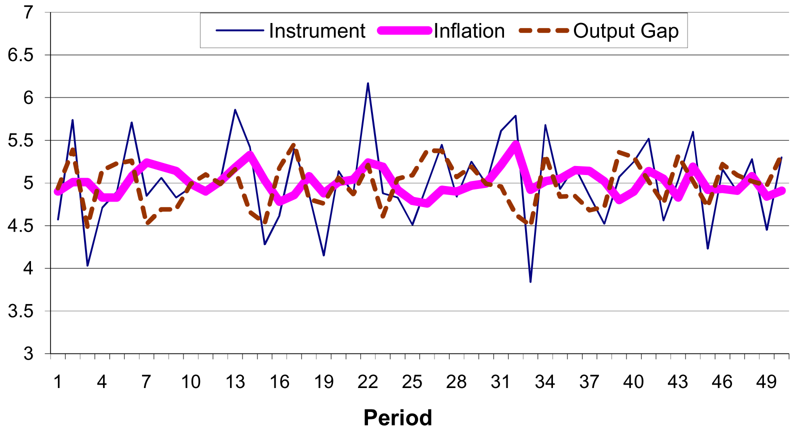

Fig. 1 Model 1 simulation using optimal rule

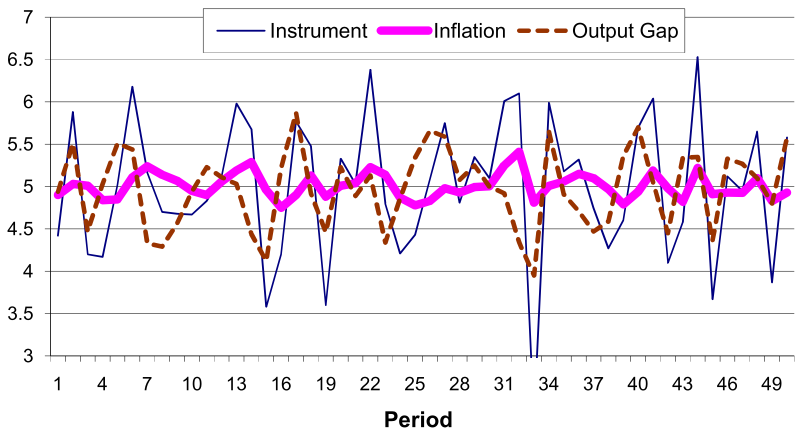

Fig. 2 Model 2 simulation using optimal rule

glean from the two figures that when optimally controlled, the time path of inflation is similar in both models.

Despite the fairly low variance in the shocks (0.15) there is still plenty of scope for monetary policy. Figures 1 and 2 show that to optimally control the system the instrument must vary between roughly 3.75 and 6.25 for Model 1, and between 3.0 and 6.5 for Model 2. Thus the decision-making problem requires the subjects to use an active strategy for controlling the economy, i.e., the problem is not a trivial one.

Sixty-eight subjects played the game with Model 1 driving the system, and of these fifty-six made nonnegative earnings (negative earnings were converted to zero earnings). The mean sum of period losses (std. dev.) for the successful subjects was 1.15 (0.98), and for the unsuccessful ones was 551.18 (1821.87). If one were to use the optimal rule total losses would be 0.57; thus the successful subjects performed very well.

Seventy subjects played the game with Model 2 driving the system, and of these fifty-eight made nonnegative earnings. The mean sum of period losses (std. dev.) for the successful subjects was 1.03 (0.89), and for the unsuccessful ones was 579.49 (1783.72). If one were to use the optimal rule total losses would be 0.57; thus, again, the successful subjects performed very well.

3.1 Aggregate results

3.1.1 Subjects who achieved control

Our data contain fifty decisions for over fifty subjects in each treatment, thus we begin our analysis with panel data estimation. Our estimation strategy takes into account two facts. First, our prior was that the decision to set the instrument in period t would be affected by the decision that was taken in period $t-1$, either because of inertia or a taste for smoothing on the part of the subjects or, as we conjecture, because of the uncertainty regarding model economy. In this case one must instrument for the lagged dependent variable on the right hand side of the equation. Second, if coefficients are heterogeneous with respect to individual subjects, imposing the restriction of homogeneity can result in severely biased estimates. We address these issues one at a time in the following analysis.

We ran a fixed effects panel regression for both models on the data from the successful subjects. For Model 1, this was fifty-six out of sixty-eight subjects (82%), and for Model 2 it was fifty-eight out of seventy subjects (83%).[13] Table 2 presents the results, located in the column labeled "Fixed Effects". We estimated an equation with the instrument as the dependent variable and the following independent variables: lagged instrument, inflation, lagged inflation, output, and lagged output. We present

[13]Eliminating the unsuccessful subjects from the analysis primarily biases the coefficient estimate on inflation upward with respect to the randomly drawn subject population. We must remove these subjects because once they lose control of the process they are constrained: the optimal strategy is not available to them any more because we limited the range of the instrument, thus our analysis is conditional on successful control. We believe this is not a problem because in real life we typically observe decisions made by people who have an understanding of the process. We also note that there was an obvious breakpoint between successful and unsuccessful performance.

Table 2 Panel regression coefficient estimates

	Fixed effects		Arrelano and Bond		Predicted	
	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2
Lagged instrument	0.41*	0.51*	0.45*	0.60*	0.00	0.00
	(0.05)	(0.05)	(0.05)	(0.05)		
Inflation	1.10*	1.00*	1.11*	1.09*	3.22	5.00
	(0.06)	(0.07)	(0.06)	(0.06)		
Lagged inflation	−0.20*	−0.53*	−0.15	−0.41*	0.00	0.53
	(0.06)	(0.07)	(0.07)	(0.07)		
Output	0.24*	0.15*	0.35*	0.37*	1.87	1.60
	(0.06)	(0.06)	(0.06)	(0.06)		
Lagged output	−0.07	0.16*	−0.12*	0.04	0.00	0.00
	(0.06)	(0.06)	(0.06)	(0.06)		
R-sq	0.41	0.49				

Notes: Model 1: 56 subjects, 49 observations per subject. Model 2: 58 subjects, 49 observations per subject.
* indicates significance at the 5% level

results from this general model; none of the conclusions we make are affected by reasonable changes to the set of independent variables.

Two coefficients are of particular importance in our analysis: the coefficients on the lagged instrument and on inflation. The regressions above indicate that our central bankers smooth interest rates: the lagged instrument affects their decision making, with positive and significant coefficients on the lagged instrument. Interestingly, our estimates of the coefficient on the lagged instrument, 0.41 for Model 1 and 0.51 for Model 2, are very much in line with the findings in Sack (1998) and Srour (2001). Sack, using quarterly data for the US between 1987 and 1997, reports the least-squares estimate of the coefficient to be 0.63. Similarly, Srour reports the estimated coefficient of 0.67; he uses quarterly data for Canada between 1984 and 1999. The estimate of the other coefficient on inflation is 1.10 for Model 1 and 1.00 for Model 2, and significantly different from zero for both. The coefficient on output is positive and significant in both models. And the coefficients on lagged inflation and output are negative and significant in most cases. The r-squared statistics for these models is in the 0.40–0.50 range, indicating that there is quite a bit of variance of decision making left to explain.

Since the lagged dependent variable might be a part of the effective strategy that subjects employ, we re-estimated the same model instrumenting for this variable using the one-stage Arellano and Bond (1991) procedure. This procedure uses all available lags of the lagged dependent variable as instruments for it. The results are presented in the second column of Table 2 labeled "Arellano and Bond".

All coefficients for both models, except lagged output in Model 2, are significant in these regressions. Once again, the lagged dependent variable appears to affect decision making in both models, though with a larger coefficient in Model 2 than in Model 1. Positive weight is placed on output for both models. Interestingly, the

Table 3 Constrained optimal coefficients

	Model 1		Model 2	
	Fixed effects	Arrelano-Bond	Fixed effects	Arrelano-Bond
Lagged instrument	0.43	0.45	0.85	0.79
Inflation	1.27	1.13	2.67	2.21
Lagged inflation	−0.13	−0.15	−0.01	−0.06
Output	2.55	2.57	2.73	2.64
Lagged output	0.23	0.11	0.63	0.37

weights on the important variables, i.e., inflation and output, appear nearly identical in both cases.

For both models and both regression specifications a 95% confidence interval for the estimated coefficient on inflation always includes 1. For Model 2 we do not reject the null hypothesis that this coefficient is 1 at the 5% level, and for Model 1 we reject at the 4% and 8% level for the fixed effects and Arellano and Bond specifications respectively. We ran a fixed effects model with all explanatory variables multiplied by a model dummy included in the right-hand side. This unreported regression fails to reject the null hypothesis that the coefficients on inflation and output are the same for both models.

In order to assess the relative success of our central bankers, we have calculated the optimal levels of single individual weights keeping the remaining coefficients at the estimated levels. To distinguish them from the theoretical optimal policy found earlier we call these coefficients "constrained optimal". These coefficients are reported in Table 3.

The weights on inflation, lagged inflation and lagged instrument seem to be chosen relatively close to the constrained optimal level in Model 1. The weight on lagged output is slightly negative whereas the constrained optimal level is 0.23. This might be a consequence of putting more negative weight on lagged inflation compared to constrained optimum. In both models the subjects put much lower weight on output compared with the constrained optimal level. Recall also that Model 2 calls for positive weight on lagged inflation. The subjects have put a negative weight instead which is in accordance with the constrained optimal levels.

Our earlier discussion suggested the existence of heterogeneity among our subjects, so we present results relaxing the restriction of homogeneity of the coefficients of the decision rule. Pesaran and Smith (1995) suggest estimating the coefficients individually and aggregating the estimates in this case, thus we ran OLS regressions on a subject-by-subject basis estimating the same model we reported in Table 2. We report the results in Figs. 3–8.

Figure 3 shows the distribution of r^2 statistics for both models. For Model 1, twenty-five computed r^2 statistics were between 0.90 and 1.0, and for Model 2 there were thirty-one. Both distributions are concentrated in the range above 0.60, the mean of both distributions is approximately 0.80, indicating a much better fit than we found with the panel regressions.

Figure 4 presents the distribution of estimated coefficients for the lagged instrument, and here we find similar results for both models with a modal estimate

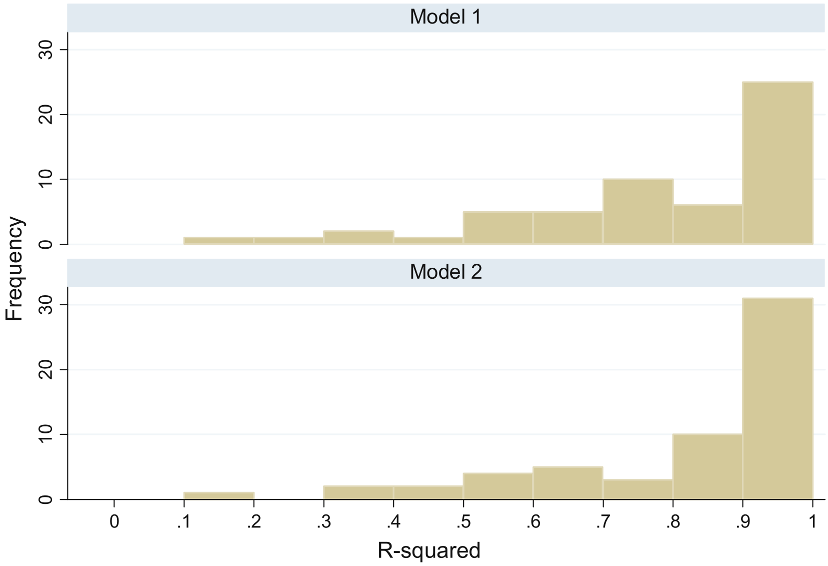

Fig. 3 Distribution of OLS R-squared

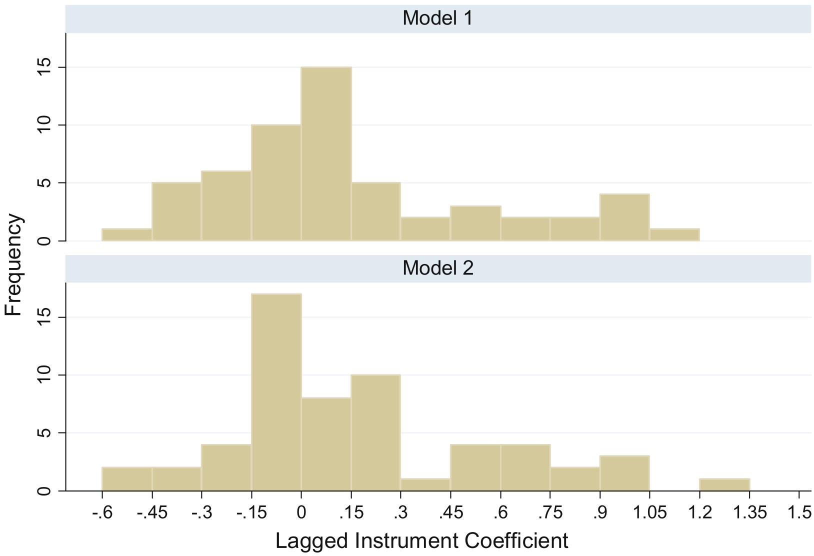

Fig. 4 Distribution of OLS estimates of lagged instrument coefficient

around 0.16, and means of 0.14 and 0.17 for Models 1 and 2, which is smaller than the panel data estimates but still positive. Most of the estimates lie between 0 and 1.

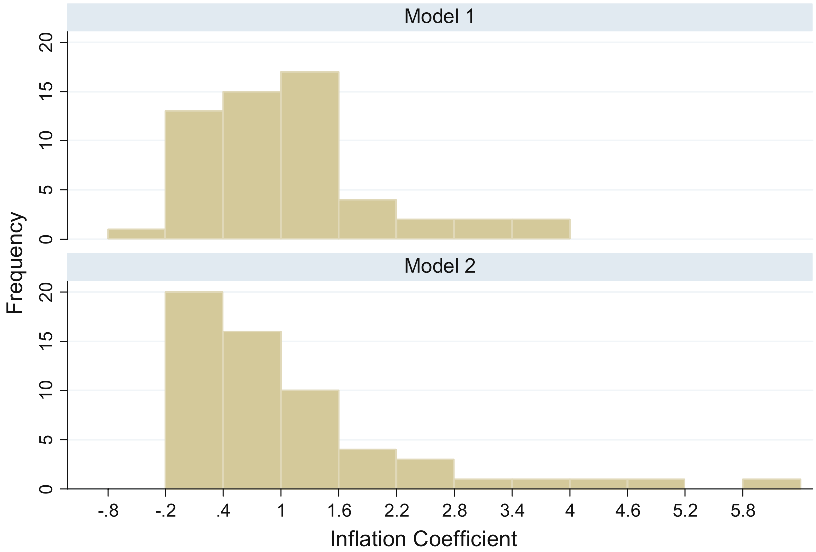

Fig. 5 Distribution of OLS estimates of inflation coefficient

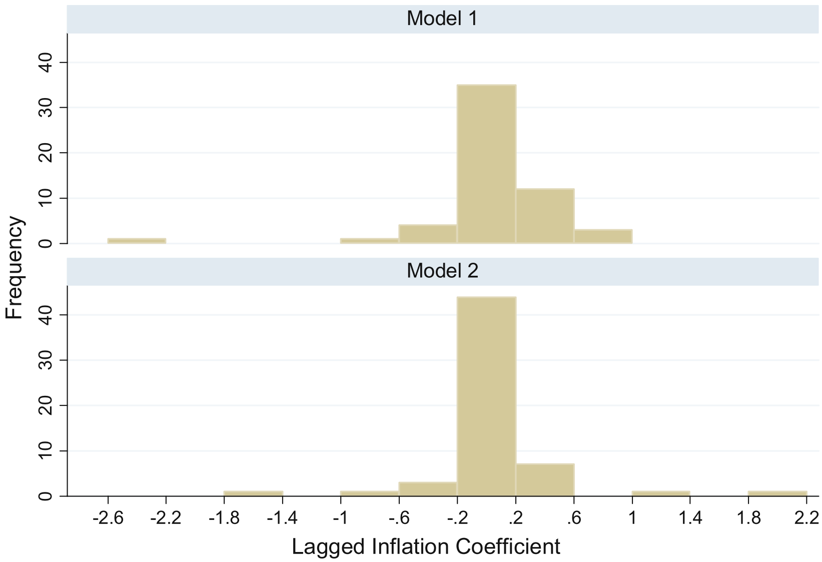

Fig. 6 Distribution of OLS estimates of lagged inflation coefficient

Figure 5 shows how individual subjects reacted to inflation. Figure 6 shows the distribution of estimated weights on the lagged inflation coefficient; the means here

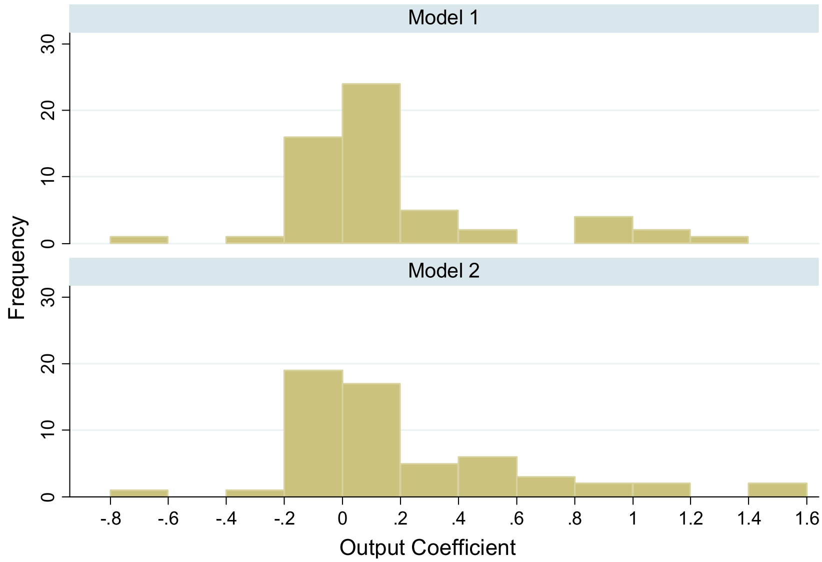

Fig. 7 Distribution of OLS estimates of output coefficient

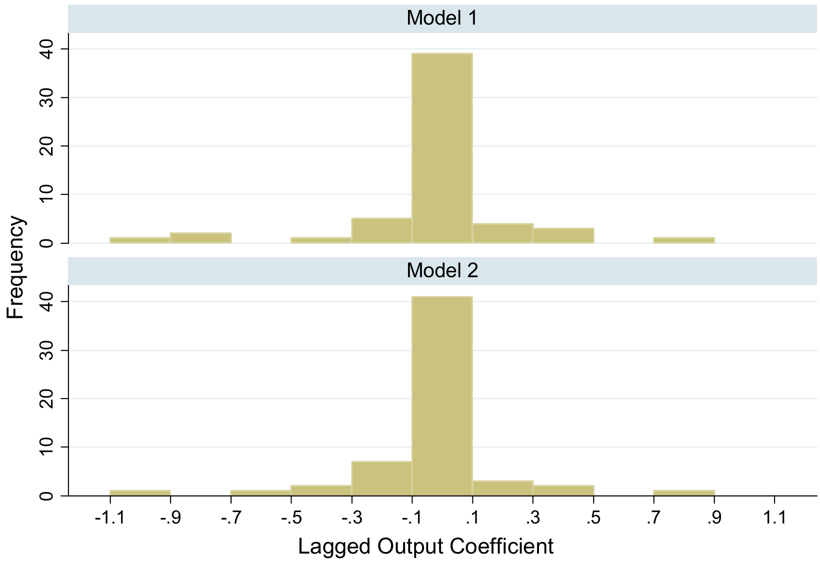

Fig. 8 Distribution of OLS estimates of lagged output coefficient

are close to zero and once again the distributions are similar. The mean weight on inflation is 1.07 and 1.11 for Models 1 and 2 respectively.

Figures 7 and 8 present estimated coefficients for output and lagged output; once again, the distributions are similar, with positive mean weight placed on output and

Table 4 Mean of OLS point estimates

	Model 1		Model 2	
	Control	No control	Control	No control
Inflation	1.11	0.59	1.16	0.17
	(0.87)	(0.85)	(1.34)	(0.64)
R-squared	0.74	0.31	0.72	0.26
	(0.23)	(0.21)	(0.28)	(0.18)
n	56	12	58	12

roughly zero mean weight placed on its first lag. The means of the estimated coefficients on output are 0.17 and 0.21 for Model 1 and Model 2, which is lower than the optimal rule coefficients.

Four things struck us with the results of the individual regressions. First, the fit of the OLS regressions was high, with averaged r-square statistics of 0.80. Second, the means of OLS estimates of the coefficients told the same general story as the estimates in the panel regressions. Third, the distributions of the estimates of the coefficients were similar across the two treatments; so similar, in fact, that a Kolmogorov-Smirnov test fails to reject the null that they are identical in every case. And fourth, these distributions of the OLS estimates look very different from what a sampling distribution of OLS estimates would look like in a sequence of repeated experiments with the same structure (i.e., asymptotically normal and resemblance to normal in a small sample). This last point bolsters the notion that the subjects are not all the same.

3.1.2 Subjects who did not achieve control

In this section we present an analysis of the decisions of subjects who did not achieve control of the economy. We do this with the caveat that whenever a subject wished to set her instrument below zero or above ten her decisions are truncated by the design of the experiment. While this design was necessary to speed up the learning process, since this occurred 43 times out of 3518 decisions in model one and 8 times out of 3636 total decisions in Model 2, it does not have a significant effect on inference of the results.

First, we ran the Arellano-Bond procedure on the twelve subjects in each model that clearly did not control the system, using the identical regression as we did previously with the subjects who did control the system. For each model, we find that the estimated weight on inflation is not significantly different than zero. Next, we ran individual OLS regressions and report the mean of the point estimates of inflation weight and r-squared of the estimated rule. We also report these means for the subjects who did achieve control.

The mean estimate of the weight on inflation is greater than one for the sample that achieved control and less than one for the sample that did not for both models. The mean r-squared is more than double with the sample that achieved control than it is for the sample that did not, again for both models. Thus the linear rule suggests that sub-optimal weight on inflation along with more noise is a characteristic of decision making that is unsuccessful.

3.2 Discussion of the results

Let us turn to stability analysis for the optimal policy rules and the estimated rules adopted by our subjects. As is shown in the on-line appendix, for both models the roots of the characteristic equation are within the unit circle, and thus the optimal rule in both models is stable.

3.2.1 Stability of non-optimal rules

Suppose the central bank follows the following rule:

$$i_t = \gamma_0 + \gamma_\pi \pi_t + \hat{\gamma}_\pi \pi_{t-1} + \gamma_x x_t + \hat{\gamma}_x x_{t-1} + \gamma_i i_{t-1}. \tag{8}$$

We will first analyze Model 1. It can be rewritten as

$$\pi_{t+1} = \alpha_1 \pi_t + \alpha_2 x_t + \alpha_3 + \epsilon_{t+1}, \tag{9}$$

$$x_{t+1} = \beta_1 x_t + \beta_2 \pi_t - \beta_2 i_t + \nu_{t+1}, \tag{10}$$

where $\alpha_3 = (1 - \alpha_1)\bar{\pi}$. Together with the policy rule it becomes a 3×3 system of difference equations. Using the criterion of stability of a 3×3 system from Farebrother (1973) and our particular parameter values $\alpha_1 = 0.5$, $\alpha_2 = 0.15$, $\beta_1 = 0.9$, $\beta_2 = 0.75$, we conclude that the system is stable if and only if $T_1 > 0$, $T_2 > 0$, $T_3 > 0$, $T_4 > 0$ (see the online-appendix for more details), where

$$\begin{aligned}
T_1 &= 2.6625 - 0.1125\gamma_\pi + 0.375\gamma_x - 1.4\gamma_i - 0.75\hat{\gamma}_x,\\
T_2 &= -0.0625 + 0.1125\gamma_\pi + 0.375\gamma_x + 0.0625\gamma_i + 0.1125\hat{\gamma}_\pi + 0.375\hat{\gamma}_x,\\
T_3 &= 2.7375 + 0.1125\gamma_\pi - 1.125\gamma_x + 2.7375\gamma_i - 0.1125\hat{\gamma}_\pi + 1.125\hat{\gamma}_x,\\
T_4 &= \left(0.6625 - 0.1125\gamma_\pi + 0.375\gamma_x - 1.4\gamma_i - 0.75\hat{\gamma}_x\right)\\
&\quad + \left[-0.3375\gamma_i + 0.1125\hat{\gamma}_\pi - 0.375\hat{\gamma}_x\right]\\
&\quad \times \left[-1.4 + 0.75\gamma_x - 0.6625\gamma_i - 0.1125\hat{\gamma}_\pi + 0.375\hat{\gamma}_x\right].
\end{aligned}$$

As a simple test of whether stability was present in our subjects' choices, we plug the estimated values of the coefficients in the central bank rule. For the fixed effects procedure we obtain the following values: $T_1 = 2.1072$, $T_2 = 0.1281$, $T_3 = 3.6574$, $T_4 = 0.3086$. Clearly, the stability conditions from Proposition 1 are satisfied.

For the Arrelano-Bond procedure, we have $T_1 = 2.1289$, $T_2 = 0.1599$, $T_3 = 3.5824$, $T_4 = 0.3100$. Again, the stability criterion is satisfied. Both procedures suggest that our subjects have learned how to steer the economy by driving the coefficients of their rules to the stability region.

Model 2 can be rewritten as

$$\pi_{t+2} = \alpha_1 \pi_{t+1} + \hat{\alpha}_1 \pi_t + \alpha_2 x_{t+1} + \alpha_3 + \epsilon_{t+2}, \tag{11}$$

$$x_{t+2} = \beta_1 x_{t+1} + \beta_2 \pi_{t+1} - \beta_2 i_{t+1} + \nu_{t+2}. \tag{12}$$

Together with the policy rule and lagged inflation this becomes a 4×4 system of difference equations. Using the criterion of stability of a 4×4 system from Farebrother (1973) and our particular parameter values $\hat{\alpha}_1 = 0.2$, $\alpha_1 = 0.3$, $\alpha_2 = 0.15$, $\beta_1 = 0.9$, $\beta_2 = 0.75$, we conclude that the system stable if and only if the following conditions hold: $T_1 > 0$, $T_2 > 0$, $T_3 > 0$, $T_4 > 0$, $T_5 > 0$, (see the on-line appendix) where

$$
\begin{aligned}
T_1 &= 1 + 0.18\gamma_i + 0.15\hat{\gamma}_x, \\
T_2 &= 3.0425 - 0.1125\gamma_\pi + 0.225\gamma_x - 1.74\gamma_i - 1.2\hat{\gamma}_x, \\
T_3 &= -0.0625 + 0.1125\gamma_\pi + 0.375\gamma_x + 0.0625\gamma_i + 0.1125\hat{\gamma}_\pi + 0.375\hat{\gamma}_x, \\
T_4 &= 1.9775 + 0.1125\gamma_\pi - 0.825\gamma_x + 1.9775\gamma_i - 0.1125\hat{\gamma}_\pi + 0.825\hat{\gamma}_x, \\
T_5 &= \left(1 + 0.18\gamma_i + 0.15\hat{\gamma}_x\right)^2\left[1.0425 - 0.1125\gamma_\pi + 0.225\gamma_x - 1.38\gamma_i - 0.9\hat{\gamma}_x\right] \\
&\quad - \left(-1.38 + 0.9\gamma_x - 1.0425\gamma_i - 0.1125\hat{\gamma}_\pi + 0.225\hat{\gamma}_x\right) \\
&\quad \times \left[\left(-1.2 + 0.75\gamma_x - \gamma_i\right)\left(-0.18\gamma_i - 0.15\hat{\gamma}_x\right)\right. \\
&\quad \left. - 0.18 + 0.150\gamma_x - 0.0425\gamma_i - 0.1125\hat{\gamma}_\pi + 0.225\hat{\gamma}_x\right].
\end{aligned}
$$

As a simple test of whether stability was present in our subjects' choices, we plug the estimated values of the coefficients in the central bank rule into the expressions above. For the fixed effects procedure, we obtain the following values: $T_1 = 1.1158$, $T_2 = 1.8844$, $T_3 = 0.1385$, $T_4 = 3.1664$, $T_5 = 0.3149$. For the Arrelano-Bond procedure, we have: $T_1 = 1.1140$, $T_2 = 1.9111$, $T_3 = 0.2053$, $T_4 = 3.0605$, $T_5 = 0.2999$. As in the case of Model 1, the obtained values for T_i indicate that all the stability conditions are satisfied. We conclude that our subjects have successfully learned how to manage the economy by following a rule with coefficients from the stability region.

We have also applied this approach to the unsuccessful subjects. For Model 1, the values of test functions are $T_1 = 1.5889$, $T_2 = -0.0099$, $T_3 = 4.8319$, $T_4 = 0.0388$. And for Model 2, these values are $T_1 = 1.0828$, $T_2 = 2.2186$, $T_3 = -0.0146$, $T_4 = 2.9153$, $T_5 = 0.1421$. As can be seen, the stability conditions are not satisfied for the failed subjects.

4 Conclusion

We infer Taylor-type linear decision rules from the observed choices of inexperienced central bankers in the experimental laboratory. Since we know the details of the economy, we bypass the typical problems in practice inferring monetary decision rules, such as unobservable inflation targets, ex-post revised macro data, and difficulties measuring an economy's potential output. In our experiments subjects set the short-term interest rate with inflation as their target.

We are able to define stabilization of the economy; we find that Taylor-type rules fit a large portion of the variance of the decisions of subjects who stabilize the economy; that subjects' behavior is consistent with interest rate smoothing; that subjects' weight on inflation is, on average, near or above 1; and that constrained to this weight, the remaining coefficient weights represent something close to a constrained optimal decision rule.

Our study is the first to use Taylor-type rules as a framework to identify interest-rate smoothing as the behavioral outcome when subjects try to achieve an inflation target. In our experiments, the central bankers, who do indeed face a type of model uncertainty, behave in accordance with earlier theoretical results that central banks smooth the interest rate when facing model uncertainty. The literature offers many explanations of interest-rate smoothing. Whether it is stability of financial markets or anchoring the public sector expectations or the credibility issue, none of these would be a source of the smoothing in our environment because the private sector behavior is captured by equations and not by subjects' decisions.

Policy rules featuring interest rate smoothing typically enhance stability of the Taylor rule and are very robust to various model uncertainty specifications (e.g., Levin et al. 1999; Orphanides and Williams 2007). As has been shown above, the optimal policy is stable and does not involve the lagged instrument. Thus, had our subjects followed the optimal policy they would not have jeopardized the economy's stability. As Table 3 reports, for the estimated levels of other coefficients the optimal weight on the lagged instrument is positive. Thus, given that the weights on other variables are at the estimated suboptimal levels, interest rate smoothing is constrained optimal. However, interest rate smoothing is not necessary for stability even if the other coefficients are fixed at the estimated levels. In other words, the system would remain stable even if the subjects put zero weight on the interest rate keeping the other weights unchanged (at the estimated levels). This can be easily seen by calculating the eigenvalues of the resulting matrix or using the stability tests discussed above. This echoes the empirical study of the US Federal Reserve policy by Rotemberg and Woodford (1997). Using a structural model of the economy, they found that the estimated policy rule could be considerably improved by using more aggressive changes of the policy instrument.[14] We conclude that our subjects find themselves in an environment of model uncertainty, and in their pursuit of attaining their goal of minimizing a loss function (inflation targeting) and stabilizing the economy (which is a necessary condition for achieving the goal), choose to smooth the interest rate.

We have taken a step in the direction of learning how inexperienced economic agents react in environments that model important issues in macroeconomics, and learning what rules are behaviorally relevant in such environments. As in the long line of experiments in IO, where inexperienced firms set prices or quantities, we begin to organize the data in the simplest of environments that we use to understand important policies.

Our study provides a foundation for exploration of a variety of issues in monetary economics. An obvious next step is to determine the parameter or set of parameters that induce different decision rules for different economies. We plan to add the exogenous variable to study behavior when exchange rates, interest rates, or output in other countries affect the economy. Indeed we can study the effect of central bank decision making in one country on another, with the ability to insert an optimal decision

[14] More generally, one could, as in Onatski and Stock (2002), evaluate the radius of changes of a particular rule that generate dynamic instability or indeterminacy.

maker as a control in one of the economies. While we ran our first study using a simple backward-looking model, it will be important to test models in which economic agents react to the central bank's decision rule. Our experiment establishes a baseline result, suggesting that within the framework of Taylor-type rules, interest-rate smoothing is useful for describing the behavior of inexperienced central bankers.

Future research may also explore other hypotheses as to why central banks smooth interest rates. It would be interesting to test whether the zero nominal interest-rate floor coupled with low inflation target would imply higher level interest-rate smoothing. And it would be an exciting experiment to test the Woodford (2003) conjecture that interest-rate smoothing may be a better policy because of its influence of the private sector expectations.

Acknowledgements We thank Alan Blinder, John Galbraith, Adrian Pagan, Daniel Parent and participants at the 2004 Economics Science Association Conference in Amsterdam, the 2004 Canadian Experimental Economics Conference in Calgary and research seminar at the New Economic School in Moscow for useful comments. We would like to thank the Editor Jordi Brandts and two anonymous referees for valuable comments and suggestions, which have significantly improved the paper. We thank Sara Constantino and Julie Héroux for research assistance. We gratefully acknowledge the Economic and Social Research Council, UK and Nuffield College, University of Oxford, UK as well as Social Sciences and Humanities Research Council (SSHRC) for financial support and the Centre for Interuniversity Research and Analysis on Organizations (CIRANO) and the Bell University Laboratory in Electronic Commerce and Experimental Economy for use of the experimental laboratory.

References

Arellano, M., & Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies*, *58*, 277–297.

Arifovic, J., & Sargent, T. (2003). Laboratory experiments with an expectational Phillips curve. In Altig, D., & Smith, B. (Eds.), *Evolution and procedures in central banking* (pp. 23–55). Cambridge: Cambridge University Press.

Ball, L. (1997). *Efficient rules for monetary policy* (Working Paper No. 5952, NBER).

Bernanke, B., & Woodford, M. (Eds.) (2005). *The inflation targeting debate*. Chicago: University of Chicago Press.

Blinder, A. (1998). *Central banking in theory and practice*. Cambridge: MIT Press.

Blinder, A., & Morgan, J. (2005). Are two heads better than one? An experimental analysis of group versus individual decision making. *Journal of Money Credit and Banking*, *37*, 789–811.

Caplin, A., & Leahy, J. (1997). The money game. *New Economy*, *4*, 26–29.

Clarida, R., Gali, J., & Gertler, M. (1999). The science of monetary policy: a new Keynesian perspective. *Journal of Economic Literature*, *37*, 1661–1707.

Cukierman, A. (1996). Why does the fed smooth interest rates. In Belongia, M. (Ed.), *Monetary policy on the 75th anniversary of the federal reserve system* (pp. 111–147). Norwell: Kluwer Academic.

Doornik, J. A. (2002). *Object-oriented matrix programming using ox* (3rd edn.). London: Timberlake Consultants Press. www.nuff.ox.ac.uk/Users/Doornik.

Farebrother, R. (1973). Simplified Samuelson condition for cubic and quartic equations. *Manchester School of Economic and Social Studies*, *41*, 396–400.

Fischbacher, U. (2007). Z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*, 171–178.

Fuhrer, J., & Moore, G. (1995). Inflation persistence. *Quarterly Journal of Economics*, *440*, 127–159.

Goodfriend, M. (1991). Interest rates and the conduct of monetary policy. In *Carnegie-Rochester conference series on public policy* (vol. 34, pp. 7–30).

Houser, D., Keane, M., & McCabe, K. (2004). Behavior in a dynamic decision problem: an analysis of experimental evidence using a Bayesian type classification algorithm. *Econometrica*, *72*, 781–822.

Judd, J., & Rudebusch, G. (1998). Taylor's rule and the Fed: 1970–1997. *Federal Reserve Bank of San Francisco Economic Review*, *3*, 3–16.

Levin, A., Wieland, V., & Williams, J. (1999). Robustness of simple monetary policy rules under model uncertainty. In Taylor, J. (Ed.), *Monetary policy rules* (pp. 263–299). Chicago: Chicago University Press.

Onatski, A., & Stock, J. H. (2002). Robust monetary policy under model uncertainty in a small model of the US economy. *Macroeconomic Dynamics*, *6*, 85–110.

Orphanides, A. (2001). Monetary policy rules based on real-time data. *American Economic Review*, *91*, 964–985.

Orphanides, A., & Williams, J. (2007). Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics*, *54*, 1406–1435.

Pesaran, M., & Smith, R. (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, *68*, 79–113.

Rotemberg, J., & Woodford, M. (1997). An optimization-based econometric framework for the evaluation of monetary policy. In *NBER macroeconomic annual* (vol. 12, pp. 297–346).

Rudebusch, G. (2001). Is the Fed too timid? Monetary policy in an uncertain world. *Review of Economics and Statistics*, *83*, 203–217.

Sack, B. (1998). Uncertainty, learning, and gradual monetary policy. In *Finance and economics discussion series*, Paper 34, Federal Reserve Board.

Srour, G. (2001). *Why do central banks smooth interest rates?* (Working Paper 2001-17). Bank of Canada.

Svensson, L. (1997a). Inflation forecast targeting: implementing and monitoring inflation targets. *European Economic Review*, *41*, 1111–1146.

Svensson, L. (1997b). Optimal inflation targets, 'conservative' central banks, and linear inflation contracts. *American Economic Review*, *87*, 98–114.

Taylor, J. (1993). Discretion vs. policy rules in practice. In *Carnegie-Rochester conference series on public policy* (vol. 39, pp. 195–214).

Taylor, J. (Ed.) (1999a). *Monetary policy rules*. Chicago: Chicago University Press.

Taylor, J. (1999b). A historical analysis of monetary policy rules. In Taylor, J. (Ed.), *Monetary policy rules* (pp. 319–341). Chicago: Chicago University Press.

Taylor, J. (1999c). The robustness and efficiency of monetary policy rules as guidelines for interest rate setting by the European Central Bank. *Journal of Monetary Economics*, *43*, 655–679.

Weymark, D. (2004). Economic structure, policy objectives, and optimal interest rate policy at low inflation rates. *North American Journal of Economics and Finance*, *15*, 25–51.

Woodford, M. (2003). Optimal interest-rate smoothing. *Review of Economic Studies*, *70*, 861–886.