

An Experimental Test Of Taylor-Type Rules With Inexperienced Central Bankers

Jim Engle-Warnick*
McGill University

Nurlan Turdaliev
University of Windsor

October 2007

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. Facing uncertainty about how the economy works, our subjects smooth the interest rate as the theoretical literature suggests they should. Our study is the first to use Taylor-type rules as a framework to identify inflation weighting 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 nos. C91, E42.

Contact Information: Engle-Warnick: Department of Economics, 855 Sherbrooke St. W., McGill University, Montreal, QC, H3A 2T7, Canada, e-mail: jim.Engle-Warnick@mcgill.ca, tel.: (514) 398-1559; fax: (514) 398-4938; Turdaliev: Department of Economics, 401 Sunset Avenue, University of Windsor, Windsor, ON, N9B 3P4, Canada, e-mail: nurlan@uwindsor.ca, tel.: (519) 253-3000, ext. 2391.

*Corresponding author.

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 (Clarida, Gali and Gertler, 1998), 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, Wieland, and Williams, 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. As mentioned in Taylor (1999c), it is important to have the inflation response coefficient γ_1 larger than 1 ('stability threshold').¹ Most of the subjects in our experiments successfully learn to get this right: the estimated γ_1 is just above 1. By contrast, the estimate for the unsuccessful subjects is less than 1. We go on to show that the primary economic incentive in the model is to achieve exactly that: despite the fact that the optimal solution calls for a higher weight on inflation, there is little increase in expected payoff for doing so.

We also find that subjects smooth the interest rate. By contrast, the optimal Taylor-type rules in the two economies do not involve a lagged instrument, i.e. they do not exhibit interest-rate smoothing. However, facing uncertainty about the model of the laboratory economy, our subjects choose to smooth.²

There are several theoretical reasons why central banks might smooth interest rates. Our results are most closely related to Sack (1998), who considered a model with two possible

slopes of the IS curve (Rudebusch (2001) also presents an uncertainty model). The actual slope of the curve is not known to the central banker. Since the intercept of the IS curve is stochastic, learning is not trivial. Thus the central banker does not have an explicit interest rate smoothing objective. Instead, it has a traditional loss function that depends on the deviations of output and inflation from their respective targets. However, the uncertainty about the macroeconomy results in the optimal policy calling for the gradual changing the interest rate. In our experiment, the subjects' payoff function does not exhibit any preference for interest rate smoothing, but there is uncertainty regarding the economic model. And we find that our central bankers smooth interest rate.³

In practice, it seems that the 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:

“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.”

Why conduct an experiment? In practice, monetary policy rules may be difficult to infer from central bank decision making for a variety of reasons. For example, Judd and Rudebusch (1988) note that the inflation target is not observable to the econometrician, causing it to be lumped in a regression constant when estimating the coefficients of the rule. Also, central bankers react to different data than are available to the econometrician: the data sets commonly used to identify decision rules contain revisions that were not available to the central bank at the time they set their instrument (Orphanides, 2001). There are also multiple methods for measuring potential output in an economy. None of these difficulties are present in laboratory data, providing the cleanest environment we can think of for studying the decision rule.

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 a weight of something close to one on inflation which is the minimum necessary to successfully manage the economy, and this weight turns out to be almost optimal.

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, Gali and Gertler, 1999). The following two aggregate equations describe the economy:

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

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

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, 1995a) 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.⁴The advantage is the great simplification of the analysis obtained by treating the private sector's expectations as an adaptive process. After some arrangements we can obtain the following model:⁵

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

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

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:

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

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 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 shown 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. 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, 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.⁷ 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.⁸

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 1999).

Subjects were able to practice the game for no pay before playing it once for pay (as in the schooling choice problem in Houser, Keane, and McCabe, 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 the same, except that the game for pay would last either 51 or 52 periods, and that they would not be told which.⁹ 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.¹⁰ 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:

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

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. 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.¹¹

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.

[insert Figures 1 and 2 here]

Figure 2 show 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 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 Individual Results

We found it rather fun to look at the time path of decision making of some of our potential future central bankers. Figures 3 - 6 present selected results from four individual subjects. Figure 3 presents a very successful subject who kept inflation almost within the narrow band of optimal control; perhaps a committee nomination is in his or her future. An important observation here is that the behavior of this subject has a considerably higher degree of interest-rate smoothing, yet achieves almost the same result as the optimal policy. This is true for almost every subject. Figure 4 shows the decisions of a subject who began to lose control in the middle of the game but seemed to be regaining it. Around period 24 this subject continued to decrease the level of the instrument even while inflation began to increase. At period 27 the subject moved the instrument above the inflation level and eventually brought it back down. The smaller inflation peak at period 42 shows the subject regaining control in a skilled manner. Maybe we should consider bringing this person in someday to troubleshoot a stormy economy.¹²

[insert Figures 3-6 here]

Figures 5 and 6 illustrate what can happen when control is lost. Figure 5 shows the decisions of a subject who makes very small changes in the instrument for 15-20 periods, then as inflation begins to decrease, decreases the instrument for too long, following with sharp increase in the instrument, then a sharp decrease, and finally another sharp increase. By contrast with Figure 4, this subject experienced wider swings over time and the future does not look promising. Figure 6 shows the decisions of another subject who made very

small changes for more than 25 periods and apparently did not know what to do with the instrument as inflation constantly decreased.

Figures 7 - 10 present selected results from four individual subjects experiencing Model 2. Figure 7 shows another example of very tight control, and Figure 8 shows a recovery. Figure 9 presents an interesting simple threshold decision rule: when inflation dropped below a threshold, the subject sharply lowered the instrument to bring inflation back up toward the target. S/he returned the instrument to the initial value and held it there, apparently waiting for another threshold to be exceeded. Figure 10 is an example of loss of control due to waiting too long to move the instrument and then not moving it enough.

[insert Figures 7 - 10 here]

3.2 Aggregate Results

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.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 results from this general model; none of the conclusions we make are affected by reasonable changes to the 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); there, using quarterly data for the US between 1987 and 1997, the author reports the least-squares estimate of the coefficient to be 0.63. The estimate of the other critical coefficient, i.e., the coefficient on inflation, is 1.10 for Model 1 and 1.00 for Model 2, and significantly different from zero for both; this coefficient should be greater than 1 to achieve control of the system. 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 labelled “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. Estimates for the coefficient on inflation are a bit greater than 1 for both models, and positive weight is placed on output for both models. Interestingly, the weights on the important variables, i.e., inflation and output, appear nearly identical in both cases. Recall also that Model 2 calls

for positive weight on lagged inflation but the subjects appear to be placing negative weight on it instead.

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 Figures 11 - 15.

Figure 11 shows the distribution of r^2 statistics for both models. Notice that the horizontal axis labels the higher of the numbers defining each bin. For example, in Figure 11, Model 1, there were 25 computed r^2 statistics that fell between 0.90 and 1.0. 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 12 presents the distribution of estimated coefficients for the lagged instrument, and here we find similar results for both models with a modal estimate 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. By far most of the estimates lie between 0 and 1.

Figure 13 shows how individual subjects reacted to inflation. A clear majority of subjects in both treatments acted as though they placed a weight greater than 1 on inflation. The mean coefficient estimates are 1.07 and 1.10 for the two treatments. Figure 14 shows the distribution of estimated weights on the lagged inflation coefficient; the means here are close to zero and once again the distributions are similar. Figures 15 and 16 present estimated coefficients for output and lagged output; once again, the distributions are similar, with positive mean weight placed on output and 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 we thought we had put together the histograms incorrectly at first. 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.

3.3 Discussion of the Results

We simulated our two models of the economy to understand the incentives faced by our inexperienced central bankers. For the first of these simulations we assumed the central bankers target inflation, i.e., they keep their eye only on container B while setting their instrument, placing zero weight on all other variables. We thought this a good starting point because it is reasonable for a subject with no experience to fixate on the container that determines her payoff, and because many subjects told us in post-experiment protocols (which were announced after the experiment and were voluntary) that they focused on Container B.

We varied the weight on inflation from 0.50 to 5.0 in increments of 0.10, averaging the sum of the period losses incurred by the central banker over 10,000 simulations of the economy. The result for Model 1 is presented in the top panel of Figure 17. The horizontal axis is the weight on inflation in the decision rule (with all other weights zero), and the vertical axis presents the expected payoff in Canadian dollars. The figure reveals that the payoff function is steep for weights on inflation less than 1.0, and then it flattens out for weights up to 3.0. The top panel of Figure 18 shows a similar situation for Model 2. Thus there is a strong incentive to get the weight just above 1.0.

Next we took the decision rule coefficients from the data, using the mean estimate for all variables except inflation, which we varied from 0 to 4.5. The results for Model 1 and Model 2 are presented in the middle panels of Figures 17 and 18. Once again the payoff function is steep roughly until the weight on inflation surpasses 1.0 and then flattens out. The figures show the locations of the means of the estimated OLS coefficients from the data, which were 1.07 for Model 1 and 1.10 for Model 2. Both means are located near where the payoff function levels off, at the top of the steep part for weights less than one.

Then we repeated the exercise varying the coefficient on output from 0 to 2.9 while holding all the weights on all other variables constant at the mean of the OLS coefficients. These results are in the lower panels of Figures 17 and 18. These panels reveal that the payoff function is flat with respect to changes in the weight given to output. All estimated expected payoffs in this range are above \$20.

How much better could the subjects have done with the optimal rule? In Model 1, the estimated expected payoff for the mean of the estimated rules is \$20.93 vs. \$21.50 for the optimal rule. In Model 2, the estimated expected payoff for the mean of the estimated rules is \$21.38 vs. \$21.9 for the optimal rule. Clearly not much is being lost for focusing attention primarily on inflation, with some interest-rate smoothing, and with sub-optimal weight on output.

It appears plausible that learning in this environment involves climbing the part of the payoff function where the weight of inflation is less than one. To investigate the hypothesis that the successful subjects climbed just to the top of the hill while the unsuccessful ones did not make it off the steep part of the hill, we ran a regression on the first 25 decisions of the unsuccessful subjects in both economies.¹⁴ We found a mean weight of 0.30 on inflation in Model 1, and a result statistically insignificant from 0 for Model 2.

While there is general agreement among macroeconomists that the weight placed on inflation must be greater than 1.0, there is less agreement as to how high above 1.0 the

weight should be. Our results illustrate that in our models a wide range of decision-rule parameters do nearly as well as the optimal rule, as long as the weight is just above 1.0. The incentive is to get the weight on inflation over 1.0, but beyond that there is not enough incentive to perfect the rule.¹⁵ Our results also illustrate the finding that within the Taylor-rule framework, interest-rate smoothing is a natural response of the inexperienced central bankers facing model uncertainty in our two computer economies.

4 Conclusion

We infer Taylor-type linear decision rules from the observed choices of inexperienced central bankers in the experimental laboratory where we know the details of the macro economy. 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. The results suggest that a linear decision rule fits a large portion of the variance of decisions of subjects who learned how to control inflation, particularly when we fit decision rules to each subject individually, and when we include a term for smoothing. 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 behave in accordance with earlier theoretical results that central banks smooth the interest rate when facing uncertainty regarding the structure of the economy. Our inexperienced central bankers, without knowledge of the underlying model, had to stabilize inflation. They successfully accomplished this task by changing the interest rate much more gradually than the optimal policy. 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. This prompts us to think of model uncertainty as the main suspect.

Our results showed that our successful bankers, on average, achieved an apparent weight of something very close to 1 on inflation, while our unsuccessful bankers achieved a weight of less than 1. We illustrated the fact that the primary behavioral incentive for bankers in this problem is to get the weight of inflation just above 1, and that there is little incentive to achieve a weight higher than that. We found that Taylor type rules, when they control for behavioral tendencies such as smoothing, are an excellent tool to get at the heterogeneity in this decision-making problem.

We do not suggest that our experiments provide evidence that central banks use Taylor-type rules, but we believe that 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. We found that inflation targeting induced behavior that was explained well by linear decision rules. And these rules exhibit interest-rate smoothing observed in the behavior of actual central bankers. Thus in an environment where the optimal rule reacts only to the current levels of inflation and output gap, our central bankers choose to smooth the interest rate. 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. And 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 eco-

nomics. 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 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 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 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.

Notes

¹Taylor 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.

²Informally, 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.

³Other 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.

⁴We 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.

⁵For more details about this kind of models, see Clarida, Gali Gertler (1999), Fuhrer and

Moore (1995a, 1995b), and Ball (1997).

⁶Our 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.

⁷This 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, Keane, and McCabe (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.

⁸These informational conditions regarding the specific economic model were similar to those in Blinder and Morgan (2005).

⁹For programming purposes we flipped a coin before the first session and used that result to run 52-period games.

¹⁰Thus, we convert the problem of minimization of loss into the problem of maximization of payoff.

¹¹Because 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.

¹²Unfortunately, anonymity requirements make it impossible to identify these two promising individuals.

¹³Eliminating 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.

¹⁴We limited this analysis to the first 25 decisions to avoid problems with constrained decisions that occurred after a loss of control.

¹⁵Houser, Keane, and McCabe (2004) made a similar finding in a dynamic decision problem mimicking the choice of schooling.

References

- Arellano, M., and 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., and Sargent T. (2003). “Laboratory Experiments with an Expectational Phillips Curve” In D. Altig and B. Smith (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. and 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. and 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. and Leahy J. (1997). “The Money Game.” *New Economy*. 4, 26–29.
- Clarida, R., Gali J., and Gertler M. (1998). “Monetary Policy Rules in Practice: Some International Evidence.” *European Economic Review*. 42, 1033–1067.
- Clarida, R., Gali J., and 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). Norewell: Kluwer Academic Publishers.

- Doornik J.A. (2002) Object-Oriented Matrix Programming Using Ox 3rd ed. London: Timberlake Consultants Press and Oxford: www.nuff.ox.ac.uk/Users/Doornik.
- Fischbacher, U. (1999). “z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter’s Manual.” Working Paper No. 21, Institute for Empirical Research in Economics, University of Zurich.
- Fuhrer, J., and Moore G. (1995a). “Inflation Persistence.” *Quarterly Journal of Economics*. 440, 127–159.
- Fuhrer, J., and Moore G. (1995b). “Monetary Policy Trade-offs and the Correlation between Nominal Interest Rates and Real Output.” *American Economic Review*. 85, 219–239.
- Goodfriend, M. (1991). “Interest Rates and the Conduct of Monetary Policy.” *Carnegie-Rochester Conference Series on Public Policy*. 34, 7-30.
- Houser, D., Keane, M., and 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., and 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., and Williams J. (1999). “Robustness of Simple Monetary Policy Rules Under Model Uncertainty.” In J. Taylor (ed.), *Monetary Policy Rules* (pp. 263-299). Chicago: Chicago University Press.
- Orphanides, A. (2001). “Monetary Policy Rules Based on Real-Time Data.” *American Economic Review*. 91, 964–985.
- Pesaran, M., and Smith R. (1995). “Estimating Long-Run Relationships from Dynamic Heterogeneous Panels.” *Journal of Econometrics*. 68, 79-113.

- Rotemberg, J., and Woodford, M. (1997). “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy.” In *NBER Macroeconomic Annual*. 12, 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.” *Finance and Economics Discussion Series Paper 34*, Federal Reserve Board.
- 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.” *Carnegie-Rochester Conference Series on Public Policy*. 39, 195–214.
- Taylor, J. (ed.), (1999a). *Monetary Policy Rules*. Chicago: Chicago University Press.
- Taylor, J. (1999b). “A Historical Analysis of Monetary Policy Rules.” In J. Taylor (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 Policy*. 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.

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
α_1^{\wedge}	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

Table 2: Panel Regression Coefficient Estimates

Pooled Regression Results						
	Fixed Effects		Arellano 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

Figure 1: Model 1 Simulation Using The Optimal Rule

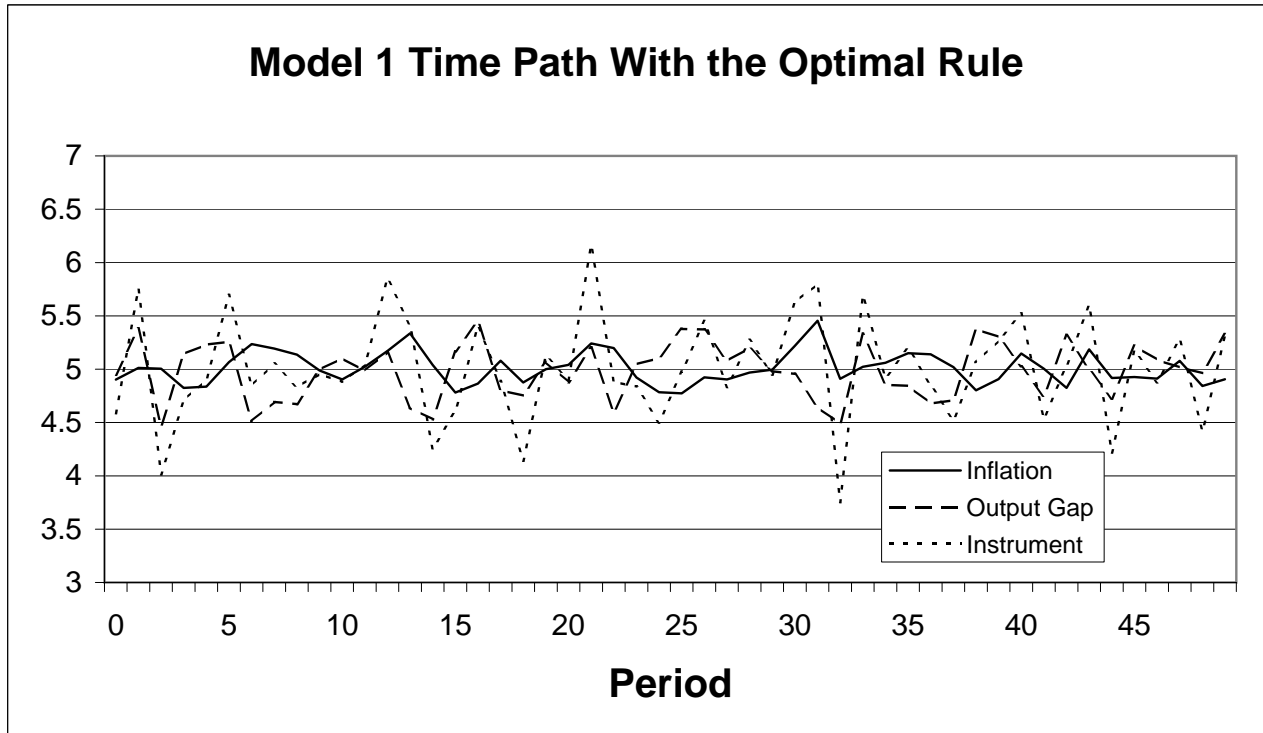


Figure 2: Model 2 Simulation Using The Optimal Rule




Figure 3: A Very Good Central Banker

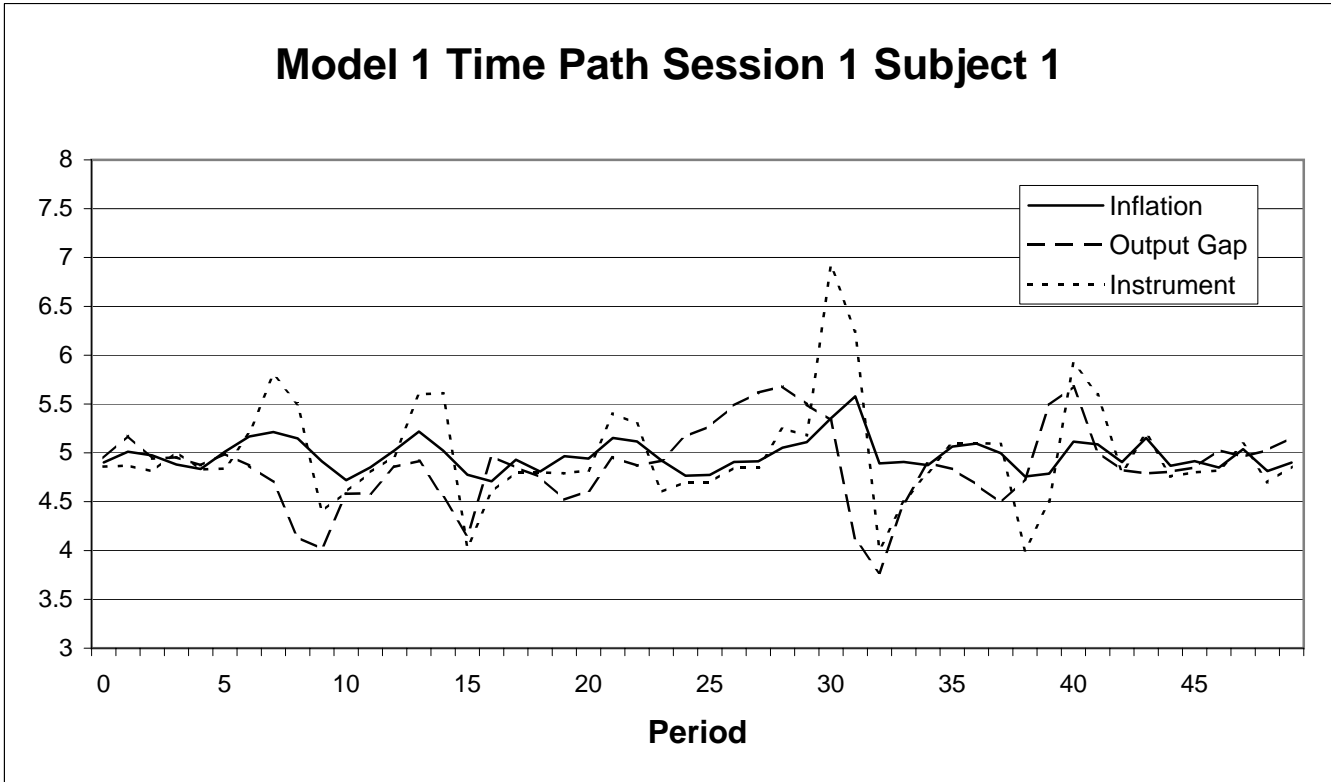


Figure 4: Increasing Variance Over Time




Figure 5: Losing Control




Figure 6:A Lost Cause




Figure 7: Another Very Good Central Banker




Figure 8: A Recovery




Figure 9: A Simple And Effective Rule

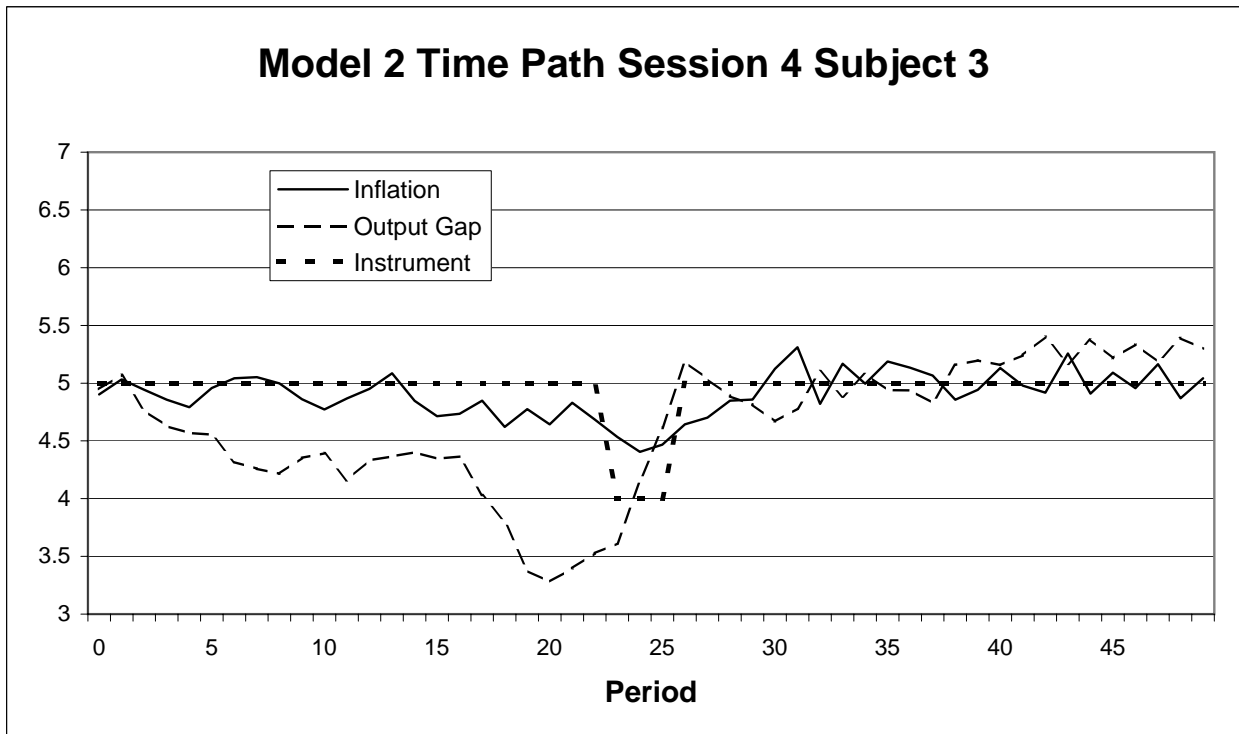


Figure 10: Another Lost Cause




Figure 11: Distribution of OLS R-squared





Figure 12: Distribution of OLS Estimates of Lagged Instrument Coefficient





Figure 13: Distribution of OLS Estimates of Inflation Coefficient





Figure 14: Distribution of OLS Estimates of Lagged Inflation Coefficient





Figure 15: Distribution of OLS Estimates of Output Coefficient





Figure 16: Distribution of OLS Estimates of Lagged Output Coefficient




Figure 17: Expected Payoffs for Model 1






Figure 18: Expected Payoffs for Model 2

