Evolution and Procedures in Central Banking

Edited by
DAVID E. ALTIG
Federal Reserve Bank of Cleveland

BRUCE D. SMITH
University of Texas

CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa

http://www.cambridge.org

First published 2003

Printed in the United Kingdom at the University Press, Cambridge

Typeface Times Roman 10/12 *System* Quark Express [AU]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication data available

ISBN 0 521 81427 8 hardback

Contents

PART I

OPERATIONAL ISSUES IN MODERN CENTRAL BANKING

1

Laboratory Experiments with an Expectational Phillips Curve

Jasmina Arifovic and Thomas J. Sargent

1. INTRODUCTION

This paper describes experiments with human subjects in an environment that provokes the time-consistency problem of Kydland and Prescott (1977). There is an expectational Phillips curve, a single policymaker, who sets inflation up to a random error term, and members of the public, who forecast the inflation rate. The policymaker knows the model. Kydland and Prescott consider a one-period model and describe how the inability to commit to an inflation policy causes the policymaker to set inflation to a Nash (that is, time-consistent) level that is higher than it would be if it could commit. With repetition (see Barro and Gordon 1983), the availability of history-dependent strategies multiplies the range of equilibrium outcomes. Some are better than the one-period, time-consistent one; others are worse.

Some commentators, including Blinder (1998) and McCallum (1995), assert that in practice, the time-consistency problem can be solved through an unspecified process that lets the monetary authority "just do it," in the terminology of an American sports shoe advertisement. Here, "it" is to choose the optimal or Ramsey target inflation rate. Although reputational macroeconomics provides no support for "just do it" as a piece of policy advice,[1] the range of outcomes predicted by that theory is big enough to rationalize such behavior. The large set of outcomes motivated us to put human subjects inside a Kydland–Prescott environment.

We paid undergraduate students to perform as policymakers and private forecasters in a repeated version of the Kydland–Prescott economy. A single policymaker repeatedly faced N forecasters, whose average forecast of inflation positioned an expectational Phillips curve.

Inspired by the theoretical literature, we ask the following questions: (1) *Emergence of Ramsey*: Is there a tendency for the optimal but time-inconsistent (Ramsey), one-period outcome to emerge as time passes within an experiment? (2) *Backsliding*: After a policymaker has nearly achieved Ramsey inflation, does inflation ever drift back toward Nash inflation? (3) *Focal points*: Are there other

[1] The theory identifies multiple systems of expectations to which the policymaker wants to conform. It provides no guidance about how to switch from one system of expectations to another.

"focal points" besides the Nash and Ramsey inflation rates? (4) *History dependence*: Is there evidence of carryover across sessions in agents' forecasts of inflation? (5) *Inferior forecasting*: Are there sometimes systematic average errors in forecasting inflation? We answer yes to the first four questions and no to the last one. The positive answer to the first question supports the "just do it" position, but it is qualified by the positive answer to the second question.

The first two questions are inspired by Barro and Gordon (1983) and Sargent (1999). Barro and Gordon describe a reputational equilibrium that can sustain repetition of the Ramsey outcome. Sargent points out that Phelps's (1967) control problem for the monetary authority under adaptive expectations for the public eventually leads the monetary authority close to Ramsey outcomes. However, Sargent also shows that repetition of the Nash equilibrium outcome is self-confirming,[2] and the "mean dynamics" of least-squares learning on the part of the government drive the system toward the self-confirming Nash equilibrium. The mean dynamics are essentially a differential version of "best response dynamics." They summarize and formalize the forces alluded to in Kydland and Prescott's heuristic sketch of an adaptive learning process that causes the government to depart from the Ramsey outcome and gradually approach the self-confirming Nash equilibrium outcome. We call this process of moving away from a Ramsey outcome, however attained, toward a Nash equilibrium "backsliding."[3]

Figure 2.1: The Nash Equilibrium and Ramsey Outcome for the Kydland–Prescott Model

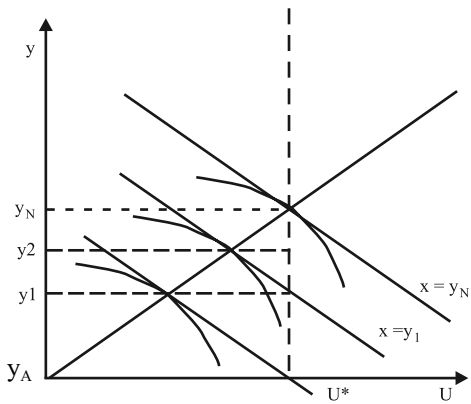

[2] A self-confirming equilibrium is a regression of unemployment on inflation that reproduces itself under a government-decision problem that takes the regression as invariant under intervention and trades inflation for unemployment. The statement in the text that the Nash equilibrium outcome is the unique, self-confirming equilibrium must be qualified because it depends on a Phillips curve that regresses unemployment on inflation. If its direction is reversed, the self-confirming equilibrium has an inflation outcome that is higher than the Nash outcome. See Sargent (1999) for details.

[3] John B. Taylor (see Solow and Taylor 1999) warns against backsliding because he believes standard time-series tests of the natural-rate hypothesis will reject it if the persistence of inflation continues to decrease, as it seems to have done in recent years in the United States.

2. THE ENVIRONMENT

Our basic model is Kydland and Prescott's. Let $(U_t, y_t, x_t, \hat{x}_t)$ denote the unemployment rate, the inflation rate, the systematic part of the inflation rate, and the public's expected rate of inflation, respectively. The policymaker sets x_t, the public sets $\hat{x}_t$, and the economy determines outcomes (y_t, U_t).

The data are generated by the natural unemployment rate model

(2.1a) $U_t = U^* - \theta\,(y_t - \hat{x}_t) + v_{1t}$

(2.1b) $y_t = x_t + v_{2t}$

(2.1c) $x_t = \hat{x}_t$,

where $\theta > 0$, $U^* > 0$, and v_t is a (2 x 1) i.i.d. Gaussian random vector with $EV_t = 0$, diagonal contemporaneous covariance matrix, and $Ev_{jt}^2 = \sigma_{vj}^2$. Here U^* is the natural rate of unemployment and $-\theta$ is the slope of an expectations-augmented Phillips curve. According to (2.1a), there is a family of Phillips curves indexed by $\hat{x}_t$. Condition (2.1b) states that the government sets inflation up to a random term, v_{2t}. Condition (2.1c) imposes rational expectations for the public and embodies the idea that private agents face a pure forecasting problem: Their payoffs vary inversely with their squared forecasting error. System (2.1) embodies the natural unemployment rate hypothesis: Surprise inflation lowers the unemployment rate, but anticipated inflation does not.

2.1. Nash and Ramsey Equilibria and Outcomes

The literature focuses on two equilibria of the one-period model. Both equilibria assume that the government knows the correct model. Called the Nash and the Ramsey equilibria, they come from different timing protocols. The Ramsey outcome is better than the Nash outcome, symptomatic of a time-inconsistency problem.

To define a Nash equilibrium, we need

DEFINITION 2.1: A government's *best response* map, $x_t = B(\hat{x}_t)$, solves the problem

(2.2) $$\min_{x_t} E\,(U_t^2 + y_t^2)$$

subject to (2.1a) and (2.1b), taking $\hat{x}_t$ as given. The best response map is

$$x_t = \frac{\theta}{\theta^2+1}\,U^* + \frac{\theta^2}{\theta^2+1}\,\hat{x}_t.$$

A Nash equilibrium incorporates a government's best response and rational expectations for the public.

Definition 2.2: A *Nash equilibrium* is a pair $(x, \hat{x})$ satisfying $x = B(\hat{x})$ and $\hat{x} = x$. A *Nash outcome* is the associated (U_t, y_t).

Definition 2.3: The *Ramsey plan* x_t solves the problem of minimizing (2.2) subject to (2.1a), (2.1b), and (2.1c). The Ramsey outcome is the associated (U_t, y_t).

A Ramsey outcome dominates a Nash outcome. The Ramsey plan is $\hat{x}_t = x_t = 0$, and the Ramsey outcome is $U_t = U^* - \theta v_{2t} + v_{1t}, y_t = v_{2t}$. The Nash equilibrium is $\hat{x}_t = x_t = \theta U,^*$ and the Nash outcome is $U_t = U^* - \theta v_{2t} + v_{1t} = \theta U^* + v_{2t}$. The addition of constraint (2.1c) on the government in the Ramsey problem makes the government achieve better outcomes by taking into account how its actions affect the public's expectations. The superiority of the Ramsey outcome reflects the value to the government of committing to a policy before the public sets expectations.

3. REPETITION

We design our experiments to implement an infinitely repeated version of the Kydland–Prescott economy. The objective of the monetary authority is to maximize

$$J = -E_0 (1-\delta) \sum_{t=0}^{\infty} \delta^t (U_t^2 + y_t^2), \ \delta \in (0,1). \tag{3.1}$$

The objective of private agents continues to be to minimize the error variance in forecasting inflation one period ahead.

Three types of theories apply to this setting.

(i) *Subgame perfection*. Reputational macroeconomics, also called the theory of credible or sustainable plans,[4] studies subgame perfect equilibria with history-dependent strategies. The theory discovers a set of equilibrium outcomes. For a large enough discount factor δ, this set includes one that repeats the Ramsey outcome forever and others that sustain worse than the one-period Nash outcome. One sensible reaction is that because it contains so many possible equilibria, the theory says little empirically.

(ii) *Adaptive expectations* (*1950s*). Suppose the government believes that the public forms expectations by Cagan–Friedman adaptive expectations:

$$\hat{x}_t = (1-\lambda) y_t + \lambda x_t \tag{3.2}$$

or $\hat{x}_t = (1-\lambda) \sum_{j=0}^{\infty} \lambda^j y_{t-j-1}$, where $\lambda \in (0,1)$.

A version of Phelps's (1967) control problem is to maximize (3.1) subject to (2.1a), (2.1b), and (3.2).

[4] See Stokey (1989) for a brief survey and Sargent (1999) for an application to the current problem.

The solution to this problem is a feedback rule,

(3.3) $x_t = f_1 + f_2 \hat{x}_t$.

With a high enough discount factor, the coefficients in (3.3) take values that make the government eventually push inflation toward the Ramsey outcome. Cho and Matsui (1995) refine this idea in the context of a broad class of expectations-formation mechanisms for the public that satisfy the same "induction hypothesis" that adaptive expectations exhibit: If sustained long enough, a constant inflation rate will eventually come to be expected by the public.[5]

(iii) *Adaptive expectations (1990s)*. Sargent (1999) shows that a self-confirming equilibrium (see Fudenberg and Levine 1993) of the Kydland–Prescott model yields the pessimistic Nash equilibrium outcome. Sims (1988), Sargent (1999), and Cho, Williams, and Sargent (2001) perturb the behavior rules of that self-confirming equilibrium by imputing to the policymaker doubts about model specification, which cause him to use a constant-gain learning algorithm. Those papers show the resulting model has both "mean dynamics," usually propelling it toward the self-confirming equilibrium, and "escape dynamics," occasionally expelling it toward the Ramsey outcome. Sample paths display recurrent, abrupt stabilizations prompted by the monetary authority's experimentation-induced discovery of an approximate natural-rate-hypothesis government, followed by gradual backsliding toward the (inferior) self-confirming equilibrium.

4. EXPERIMENTS

4.1. Design

A group of $N + 1$ students composes the economy; we set N equal to 3, 4, or 5. The first N students form the public. Their decision is to forecast the inflation rate for each period of the experiment. Call agent i's forecast x_{it} and let x_t be the average of the citizens' forecasts. Citizens receive payoffs that rise as their session-average squared forecast errors fall. Agent i's payoff at the end of time period t is given by

$$-.5\,(y_t - x_{i,t})^2.$$

Student $N + 1$, chosen at random at the beginning of an experiment, is the policymaker. Each period, student $N + 1$ sets a target inflation rate, x_t. A random number generator sets v_{2t} and the actual inflation rate equals $y_t = x_t + v_{2t}$.

[5] Cho and Matsui (1999) study a version of the repeated model with alternating choices by the government and the public. They find that, depending on relative discount factors, the one-period Nash outcome is excluded in an equilibrium outcome, and a narrow range of outcomes near Ramsey can be expected under some parameter settings.

Unemployment is then generated by the Phillips curve (2.1a). Student $N + 1$'s payoff varies inversely with the session average of $U_t^2 + y_t^2$ and is given by

$$-.5\,(U_t^2 + y_t^2).$$

The same student remains the policymaker throughout all sessions within a single experiment. Sessions within an experiment are separated by a stopping time.

4.2. Knowledge

The policymaker knows the true Phillips curve (2.1); the existence of private agents who are trying to forecast its action; and the histories of outcomes (y_t, U_t) in the current experiment up to the current time. The private forecasters know the history of inflation and unemployment, including prior sessions of the current experiment. At the beginning of the economy, there is no history. The private forecasters do not know the structure of the economy. They know that a policymaker sets inflation up to a random term.[6]

4.3. Physical Details

Subjects sat at computer terminals and were isolated from one another. They received written instructions at the beginning of each experiment. Appendixes A and B reproduce the instructions. All experiments were conducted at the microcomputer lab of Simon Fraser University, in Burnaby, Canada. Subjects were undergraduate economics majors at Simon Fraser University. They were recruited for two-hour experiments but were not told in advance how many sessions would be played during each experiment. No subject was used in more than one experiment.

We conducted a total of 12 experiments, three in April 1998 and nine between February and April 1999.

4.4. Stopping Rule

We followed Duffy and Ochs (1999) and Marimon, McGrattan, and Sargent (1990) in using a random stopping rule to implement an infinite horizon and to discount future payoffs with the discount factor $\delta \in (0,1)$. At the end of each period, the computer program drew a random number from a uniform distribution over [0, 1]. If this random number was less than δ, the experimental session would continue for one more period. If the number was greater than δ, the session was terminated. An upper bound on the duration of an individual session was set at 100 time periods.

6 The experiments implement the environment described by Kydland and Prescott (1977), in which the government knows the model. Our assumptions about what the government and private forecasters know differ from those in Sargent (1999) and Cho, Williams, and Sargent (2002), where the private agents know the government's rule for setting the predictable part of inflation, and the government does not know the true Phillips curve model but estimates a nonexpectational Phillips curve.

4.5. Earnings

Subjects received a $10 payment (Canadian funds) for completing a two-hour experiment. They also could earn an additional $10 prize,[7] determined in the following way: At the end of each experimental time period, the number of *period points* was calculated by adding 100 points to the subject's payoff. If this number was less than 0, it was truncated to 0. Then the number of *total points* was calculated by adding all period points earned in a session. Finally, the number of *maximum points* was calculated as the product of 100 and the number of session periods. At the end of a session, a probability of winning the prize, π_{win}, was computed as the ratio between the total points and the maximum points.

Once an experiment was over, the computer program chose one of the sessions at random and chose a number, *rand*, from a uniform distribution over [0, 1]. If π_{win} of the selected session was greater than *rand*, the subject earned an additional $10. The parameter values used in the experiments were $U^* = 5$, $\theta = 1$, and the discount parameter was $\delta = 0.98$. Two sets of values of the noise standard deviation σ were used, $\sigma \equiv \sigma_1 = \sigma_2 = 0.3$. In addition to the setting of σ, an information variable (yes or no) recorded whether the policymaker had been told the value of $\hat{x}_t$ from the previous period.[8]

Each experiment was labeled an "economy" and consisted of a set of sessions with the same policymaker and group of forecasters. Each economy had several sessions. Table 4.1 summarizes the treatment variables across economies.[9]

Table 4.1: Design of Experiments

Experiment	Sessions	Information	σ	N
1	3	*	.03	4
2	2	**	.03	4
3	3	***	.3	5
4	2	yes	.3	3
5	2	yes	.3	4
6	9	yes	.3	4
7	6	yes	.3	4
8	9	yes	.3	4
9	4	yes	.3	4
10	2	yes	.3	4
11	9	yes	.3	4
12	9	yes	.3	4

7 We used a version of the Roth–Malouf (1979) binary lottery to determine actual cash payments with the intention to control for subjects' differing attitudes toward risk.

8 We used two alternative scales for the payoffs for the forecasters. For experiments 1–8, we used $-.5(y_t - \hat{x}_{it})^2$, while for experiments 9–12, we used $-5\,(y_t - x_{it})^2$. The second scale was introduced to increase the weight of poor forecasts in the calculation of π_{win}.

9 In table 4.1, (*) denotes (no, yes, yes), (**) denotes (no, yes), and (***) denotes (no, yes, yes) in successive sessions.

5. OUTCOMES

Tables 5.1, 5.2, and 5.3 and figures C1–17 describe the outcomes. Each economy corresponds to one set of $N + 1$ students. Figures C14–17 contain evidence about the heterogeneity of the citizens' expectations of inflation. Each economy contains several sessions, determined by the realization of a random variable that terminated the session. The panels in figures C1–12 correspond to different sessions with the same group of students.

Table 5.1 reports the means and standard deviations of $\hat{x}_t$, x_t, y_t, U_t, $-.5\,(U_t^2 + y_t^2)$ across all sessions for each group. For the parameter values $U^* = 5$ and $\theta = 1$, the population values for these variables at the Nash equilibrium are 5, 5, 5, 5, and –25. For the Ramsey outcome, the values are 0, 0, 0, 5, and –12.5.

5.1. Patterns

Table 5.2 summarizes the patterns in figures C1–17. The column labels represent the following: "Ramsey" indicates the policymaker pushes the system to Ramsey at least for a substantial length of time (see figures C1 and C2 for economies 1 and 2). "Backsliding" indicates a resurgence of inflation after having attained Ramsey (see figures C3 and C6). "Other focal" indicates sustained inflation at values distinct from the Ramsey or Nash inflation (see figure C9). "Experimentation" indicates the presence of episodes in which the monetary authority seems to be engaging in purposeful experimentation. "Rank" denotes the rank order of the experiments in terms of the economywide average payoff for the monetary authority. An "x" signifies strong evidence for the pattern in question, a "y" signifies weaker evidence, and a blank signifies no evidence. Table 5.3 reports the results of regressions of inflation and the government payoff, respectively, on a constant and a dummy that takes value 0 in the first half of an experiment and 1 in the second half, where the second half is defined as the last $N/2$ sessions if N is even and $N - 1/2$ sessions if N is odd. The table reports regression coefficients with standard errors in parenthesis; an asterisk denotes statistical significance at the 5 percent level. We summarize the main features of the results as follows.

- Figures C1–12 indicate that, on average, the public's forecasts of inflation are good and do not contain systematic forecast errors.
- In 9 of the 12 experiments, the policymaker pushes inflation near the Ramsey value for many periods.
- Backsliding occurs in 4 of 12 economies.
- Table 5.3 indicates that inflation falls and government payoffs rise during the second half of 10 of the 12 experiments; the decrease in inflation is statistically significant in 9 of the 12 experiments.
- The policymaker experiments in 3 of 12 economies.

Table 5.1: Means and Standard Deviations of Outcomes

Economy	x	$\hat{x}$	y	U	Gov. payoff
Nash	5	5	5	5	–25
Ramsey	0	0	0	5	–12.5
1	4.1173 (1.5267)	4.1497 (1.4923)	4.1125 (1.5298)	5.0381 (.4671)	–22.4196 (5.2823)
2	1.4937 (2.2521)	1.5047 (2.2286)	1.4888 (2.2522)	5.0183 (0.8135)	–16.5486 (7.2296)
3	1.1266 (1.1115)	1.1455 (1.0726)	1.1162 (1.1347)	5.0263 (0.5334)	–14.0370 (3.1575)
4	1.3326 (0.7794)	1.4218 (0.8094)	1.2930 (0.8360)	5.1438 (0.5383)	–14.5550 (2.8898)
5	2.0143 (1.7884)	2.2536 (1.7682)	1.9998 (1.8025)	5.2495 (1.1115)	–18.0040 (7.5711)
6	1.9196 (2.8144)	2.0600 (2.3279)	1.9086 (2.8319)	5.1636 (2.1278)	–21.4142 (26.5034)
7	1.3561 (1.1482)	1.4444 (1.1892)	1.3080 (1.1962)	5.0956 (0.6071)	–14.7334 (3.8102)
8	0.7879 (0.7897)	0.8354 (0.9031)	0.7582 (0.8551)	5.0545 (0.4979)	–13.5492 (3.0613)
9	5.8802 (1.9699)	5.8129 (1.7939)	5.8274 (1.9725)	4.9490 (1.1919)	–31.8680 (8.7549)
10	2.4640 (2.4087)	2.5443 (2.0490)	2.4158 (2.4543)	5.1006 (1.9718)	–20.8438 (11.6304)
11	3.6396 (0.7379)	3.6664 (0.7217)	3.6158 (0.7873)	5.0216 (0.7706)	–19.7498 (4.6579)
12	2.6957 (1.7212)	2.7048 (1.1878)	2.6659 (1.7263)	5.0161 (1.5879)	–18.8765 (15.8123)

Table 5.2: Patterns of Results

Economy	Ramsey	Backsliding	Other focal	Experimentation	Rank
1	x				11
2	x				5
3	x	x			2
4	x				3
5	x	x		x	6
6	x	x	y	x	9
7	x				4
8	x				1
9			y		12
10	x	x		x	10
11			x		8
12			x		7

Table 5.3: Inflation and Government Payoff on Second-Half Dummy

Economy	Inf. incpt	Inf. dummy	Gov. incpt	Gov. dummy
1	4.5637	–2.8665*	–23.7411	8.4138*
	(0.0882)	(0.2226)	(0.3396)	(0.8571)
2	3.8583	–3.8385*	–22.3969	9.4744*
	(0.1596)	(0.2031)	(0.7084)	(0.9017)
3	1.3255	–1.0571*	–14.1841	0.7430
	(0.0829)	(0.8164)	(0.7084)	(0.9017)
4	1.6127	–0.9593*	–15.1001	1.6353*
	(0.0819)	(0.1419)	(0.3251)	(0.5631)
5	3.1016	–1.9891*	–21.4088	6.1463*
	(0.1917)	(0.2576)	(0.8823)	(1.1854)
6	1.9549	–0.2388	–21.9732	2.8805
	(0.1356)	(0.3079)	(1.2692)	(2.8811)
7	1.6024	–0.9722*	–15.0774	1.1365*
	(0.0840)	(0.0931)	(0.2088)	(0.3323)
8	0.8524	–0.2388*	–13.9437	0.9993*
	(0.0585)	(0.0931)	(0.2088)	(0.5195)
9	5.4664	1.0468	–30.7858	–3.1381
	(0.1658)	(0.2824)	(0.7498)	(1.2769)
10	2.8456	–1.2428*	–21.8264	2.8410
	(0.2562)	(0.4357)	(1.2431)	(2.1137)
11	3.4265	0.3872	–18.8962	–1.7463
	(0.0533)	(0.0763)	(0.3199)	(0.4576)
12	3.1043	–1.1103*	–19.0925	0.5472
	(0.1132)	(0.1802)	(1.0925)	(1.7388)

- Economy 9 has a bad or indifferent policymaker. He attains an average payoff level worse than that associated with the Nash outcome—the only policymaker to fall short of the Nash outcome.
- Most of the transitions from Nash to Ramsey are smooth. Few (if any) have the drama of the Volcker-like rapid disinflations produced by the escape-route dynamics of Cho, Williams, and Sargent (2002) and Sargent (1999). Depending on parameter values, they could resemble a pattern predicted by Phelps (1967) and Cho and Matsui (1995). However, the stabilizations are too slow to be explained in this way, at least if policymakers are assumed to know the rate at which the public is adapting its expectations.
- Heterogeneity of expectations across citizens is largest at the beginning of an experiment. It also tends to grow at the start of a new session within an experiment.

6. ADAPTIVE EXPECTATIONS

To check whether the results confirm the predictions of the Phelps (1967) problem, we estimated the parameter λ in the adaptive-expectations model (3.2). We estimated the model for each individual within an experiment, pooling across sessions,[10] and for the average of households within an experiment, pooling across sessions.[11] For econometric reasons, we wrote the model in the form

$$(6.1) \qquad \hat{x}_t = (1 - \lambda_i) \sum_{j=0}^{t-1} \lambda_i^j y_{t-j} + \eta_i \lambda_i^t + u_{it},$$

where u_{it} is a random disturbance with mean zero that is orthogonal to y_{t-1-j} for $j = 0, ..., t - 1$, and η is the systematic part of the initial condition (see Klein 1958). We estimated (6.1) by maximum likelihood, assuming a Gaussian distribution for u_{it}. For each individual, we pooled across sessions, estimating a common λ_i but a different, session-specific η_i for each session. For the average of forecasts across individuals, x_t, we proceeded in a similar way, estimating a common λ across sessions as well as session-specific η's.

Table 6.1 shows the estimates of λ,[12] most of which are below .5, indicating that most citizens formed forecasts by heavily overweighting the recent past. In the next section, we will study whether policymakers can be viewed as solving a Phelps problem in light of this rapid adjustment.

6.1. Adaptive Expectations with Heteroskedasticity

Tables 6.2 and 6.3 summarize some of the results of re-estimating the adaptive-expectations model (6.1) by maximum likelihood while allowing the variance of the disturbance u_{it} to vary across the two halves of an experiment, defined the same way as table 5.3.[13] Table 6.3 reports estimates of the variances across the two halves, denoted σ_2^1, σ_2^2, respectively, as well as an estimate σ_2^2 that imposed homoskedasticity across the two halves. An asterisk by σ_2^2 denotes the difference across the two halves is statistically significant at the 5 percent level, according to a Chi-square test. In most experiments and for most of the private agents, the variance of u_{it} fell across the two halves of an experiment.

6.2. Phelps Problem

In the row labeled L.S., table 6.4 records least-squares estimates of the government's rule (3.3). In the row labeled Phelps, the table also reports the rule that solves the Phelps problem for $\delta = .98$ and the value of λ from table 6.1 for the

[10] Thus, there is one λ_i for each subject.
[11] Here there is one λ for each experiment for each individual.
[12] In experiment 3, there is a fifth private agent. His/her estimate of λ_i is .2303 (.0314) with an R^2 of .9938.
[13] There is a fifth agent in experiment 3, with estimated $\lambda = .2310$ (.0276). For the fifth agent, we estimated $\sigma_2 = .0168$, $\sigma_1^2 = .0199$, $\sigma_2^2 = .0097$. The difference in disturbance variances across halves is not statistically significant at the .05 level.

Table 6.1: Estimates of λ_i in (6.1)

Exp.		Agent 1	Agent 2	Agent 3	Agent 4	Average
1	λ	0.1395	0.0942	0.3136	0.1618	0.1896
	s.e.	(0.0350)	(0.0828)	(0.0549)	(0.0436)	(0.0322)
	R^2	0.9983	0.9919	0.9952	0.9972	0.9986
2	λ	0.1698	0.1366	0.2501	0.0007	0.1950
	s.e.	(0.0915)	(0.0885)	(0.0506)	(0.0015)	(0.0382)
	R^2	0.9475	0.9736	0.9656	0.9692	0.9912
3	λ	0.3278	0.4007	0.3363	0.4627	0.3556
	s.e.	(0.0649)	(0.0452)	(0.0381)	(.0359)	(0.0278)
	R^2	0.9737	0.9809	0.9897	0.9862	0.9938
4	λ	0.7345	0.3849	0.2635		0.4126
	s.e.	(0.0805)	(0.0641)	(0.0638)		(0.0547)
	R^2	0.9755	0.9852	0.9820		0.9893
5	λ	0.5059	0.3644	0.2605	0.7918	0.5006
	s.e.	(0.0493)	(0.0609)	(0.0605)	(0.0343)	(0.0360)
	R^2	0.9569	0.9498	0.9539	0.9092	0.9846
6	λ	0.8413	0.7829	0.7059	0.6335	0.7452
	s.e.	(0.0232)	(0.0635)	(0.0147)	(0.0292)	(0.0137)
	R^2	0.7844	0.5853	0.8761	0.8569	0.9461
7	λ	0.2585	0.3362	0.7562	0.3124	0.4160
	s.e.	(0.0409)	(0.0295)	(0.0368)	(0.0505)	(0.0746)
	R^2	0.9840	0.9929	0.6209	0.9801	0.9692
8	λ	0.4893	0.4200	0.2788	0.3632	0.3935
	s.e.	(0.0370)	(0.0329)	(0.0014)	(0.0048)	(0.0236)
	R^2	0.9544	0.9691	0.9696	0.9769	0.9872
9	λ	0.5649	0.1233	0.2662	0.2073	0.3392
	s.e.	(0.0214)	(0.0730)	(0.0405)	(0.0503)	(0.0201)
	R^2	0.9960	0.9907	0.9940	0.9875	0.9983
10	λ	0.1300	0.2877	0.4176	0.6609	0.3800
	s.e.	(0.0476)	(0.1145)	(0.0322)	(0.0438)	(0.0442)
	R^2	0.9368	0.4668	0.9667	0.9151	0.9393
11	λ	0.4796	0.4856	0.5322	0.4378	0.5109
	s.e.	(0.0244)	(0.0422)	(0.1162)	(0.0356)	(0.0367)
	R^2	0.9966	0.9861	0.8596	0.9915	0.9888
12	λ	0.7083	0.0663	0.7136	0.4836	0.4776
	s.e.	(0.0366)	(0.0222)	(0.0476)	(0.0410)	(0.0264)
	R^2	0.8338	0.9533	0.9249	0.9255	0.9708

Table 6.2: Estimates of λ_i with Heteroskedasticity

Economy		Agent 1	Agent 2	Agent 3	Agent 4	Average
1	λ	0.1107	0.0100	0.3409	0.0698	0.1374
	s.e.	(0.0229)	(0.0073)	(0.0473)	(0.0433)	(0.0203)
2	λ	0.1842	0.1645	0.3222	0.0364	0.2176
	s.e.	(0.0522)	(0.0454)	(0.0573)	(0.0602)	(0.0373)
3	λ	0.2675	0.4102	0.3312	0.4499	0.3616
	s.e.	(0.0508)	(0.0444)	(0.0408)	(0.0344)	(0.0277)
4	λ	0.5592	0.3642	0.2519		0.3727
	s.e.	(0.1141)	(0.0564)	(0.0586)		(0.0548)
5	λ	0.5582	0.3690	0.2750	0.6021	0.5030
	s.e.	(0.0436)	(0.0598)	(0.0671)	(0.0657)	(0.0307)
6	λ	0.8220	0.0009	0.7153	0.7187	0.7571
	s.e.	(0.0183)	(0.0815)	(0.0156)	(0.0347)	(0.0154)
7	λ	0.2075	0.3321	0.3259	0.3076	0.3288
	s.e.	(0.0288)	(0.0274)	(0.0728)	(0.0468)	(0.0397)
8	λ	0.5618	0.4133	0.2981	0.3466	0.4087
	s.e.	(0.0264)	(0.0327)	(0.0347)	(0.0317)	(0.0207)
9	λ	0.5625	0.2461	0.3094	0.2021	0.3574
	s.e.	(0.0224)	(0.0509)	(0.0425)	(0.0369)	(0.0210)
10	λ	0.1639	0.3319	0.3958	0.6542	0.3866
	s.e.	(0.0478)	(0.0744)	(0.0327)	(0.0438)	(0.0458)
11	λ	0.4917	0.5969	0.4790	0.3748	0.5323
	s.e.	(0.0234)	(0.0242)	(0.0307)	(0.0367)	(0.0203
12	λ	0.8196	0.0169	0.7855	0.4647	0.4696
	s.e.	(0.0175)	(0.0160)	(0.0307)	(0.0391)	(0.0253)

averaged-across-individuals values of $\hat{x}_t$. The least-squares estimates of (3.3) show the policymakers seem to have adjusted inflation downward too slowly relative to the solution to the Phelps problem. In particular, the least-squares values of f_2 are always substantially larger than those associated with the optimal rule from the Phelps problem. If policymakers are to be interpreted as solving a Phelps problem, then they must be regarded as acting as though they think members of the public adjust much more slowly (have higher λ) than they apparently do.

Table 6.3: Restricted vs. Unrestricted MLE

Economy		Agent 1	Agent 2	Agent 3	Agent 4	Average
1	σ^2	0.0321	0.1544	0.0906	0.0541	0.0265
	σ_1^2	0.0367	0.1816	0.1015	0.0622	0.0312
	σ_2^2	0.0085*	0.0165*	0.0330*	0.0182*	0.0043*
2	σ^2	0.3784	0.1903	0.2476	0.2208	0.0635
	σ_1^2	1.1004	0.5318	0.6957	0.6105	0.2080
	σ_2^2	0.0034*	0.0042*	0.0057*	0.0043*	0.0017*
3	σ^2	0.0659	0.0477	0.0257	0.0345	0.0154
	σ_1^2	0.0791	0.0499	0.0252	0.0303	0.0161
	σ_2^2	0.0141*	0.0389	0.0278	0.0515	0.0129
4	σ^2	0.0574	0.0348	0.0421		0.0250
	σ_1^2	0.0713	0.0410	0.0492		0.0298
	σ_2^2	0.0343	0.0226	0.0283		0.0160
5	σ^2	0.3089	0.3603	0.3304	0.6516	0.1102
	σ_1^2	0.4310	0.4824	0.3543	1.3211	0.1674
	σ_2^2	0.2168*	0.2635	0.3118	0.2209*	0.0649*
6	σ^2	2.4741	4.7589	1.4214	1.6426	0.6182
	σ_1^2	2.6390	5.8032	1.6929	1.9862	0.7254
	σ_2^2	1.8053	0.4827*	0.3028*	0.3424*	0.1787*
7	σ^2	0.0492	0.0198	1.1652	0.0613	0.0945
	σ_1^2	0.0666	0.0220	0.18567	0.0577	0.1396
	σ_2^2	0.0106*	0.0269	0.0238	0.0693	0.0135*
8	σ^2	0.0574	0.0384	0.0380	0.0278	0.0153
	σ_1^2	0.0879	0.0440	0.0482	0.0332	0.0209
	σ_2^2	0.0131*	0.0298	0.0228*	0.0195*	0.0070*
9	σ^2	0.1474	0.3448	0.2236	0.4676	0.0641
	σ_1^2	0.1823	0.5061	0.3063	0.6947	0.0901
	σ_2^2	0.0816*	0.0534*	0.0710*	0.0395*	0.0157*
10	σ^2	0.7407	6.2520	0.3909	0.9958	0.7112
	σ_1^2	0.5203	8.9559	0.2525	0.8240	0.8237
	σ_2^2	1.1613	1.2152*	0.6536	1.3176	0.5013
11	σ^2	0.0441	0.1870	1.8837	0.1137	0.1499
	σ_1^2	0.0377	0.3194	3.6244	0.0818	0.2655
	σ_2^2	0.0509	0.0567*	0.0740*	0.1487	0.0298*
12	σ^2	1.6382	0.4612	0.7400	0.7296	0.2879
	σ_1^2	2.5643	0.6409	0.9146	0.5356	0.2661
	σ_2^2	0.3454*	0.2042*	0.4928*	1.0261	0.3213

Table 6.4: Estimates of Phelps Rule (3.3)

Experiment		f_1	f_2	n	R^2	λ
1	*L.S.*	0.0515	0.9798	191	0.9172	0.1896
	s.e.	(0.0944)	(0.0214)			
	Phelps	0.0779	0.3655			
2	*L.S.*	0.0672	0.9480	162	0.8800	0.1950
	s.e.	(0.0743)	(0.0277)			
	Phelps	0.0785	0.3649			
3	*L.S.*	0.0238	0.9627	202	0.8630	0.3556
	s.e.	(0.0425)	(0.0271)			
	Phelps	0.0997	0.3504			
4	*L.S.*	0.0349	0.9127	111	0.8984	0.4126
	s.e.	(0.0480)	(0.0294)			
	Phelps	0.1099	0.3456			
5	*L.S.*	0.1190	0.8410	139	0.6914	0.5006
	s.e.	(0.1373)	(0.0480)			
	Phelps	0.1298	0.3386			
6	*L.S.*	0.2215	0.8243	541	0.4649	0.7452
	s.e.	(0.1183)	(0.0381)			
	Phelps	0.2509	0.3233			
7	*L.S.*	0.0488	0.9051	251	0.8787	0.4160
	s.e.	(0.0398)	(0.0213)			
	Phelps	0.1105	0.3453			
8	*L.S.*	0.0771	0.8508	347	0.9467	0.3935
	s.e.	(0.0134)	(0.0109)			
	Phelps	0.1063	0.3472			
9	*L.S.*	0.5250	0.9213	203	0.7038	0.3392
	s.e.	(0.2564)	(0.0422)			
	Phelps	0.0971	0.3519			
10	*L.S.*	0.5248	0.7622	133	0.4204	0.3800
	s.e.	(0.2551)	(0.0782)			
	Phelps	0.1038	0.3484			
11	*L.S.*	1.1289	0.6848	401	0.4486	0.5109
	s.e.	(0.1420)	(0.0380)			
	Phelps	0.1326	0.3378			
12	*L.S.*	0.8612	0.6782	347	0.2191	0.4776
	s.e.	(0.2036)	(0.0689)			
	Phelps	0.1240	0.3404			

7. SESSION DEPENDENCE

Figures C1–12 display visual evidence of what we call "session dependence," a tendency of the monetary authority to set the systematic part of inflation equal to its value at the end of the preceding session within an experiment. A regression of the beginning-of-session setting of x against the previous session's last setting of x, pooled across sessions and experiments, shows there is some such tendency, but it is weak:

$$x_{1(j)} = \underset{(.70)}{1.71} + \underset{(.25)}{.67} x_{T(j-1)},$$

where standard errors are in parentheses, $R^2 = .14$, $x_{1(j)}$ is the first-period setting of x within session $j \geq 2$, and $x_{T(j-1)}$ is the last-period setting of x within session $j - 1$.

8. DISPERSION

Figure C13 displays sample variances of individual forecasts of inflation, $\hat{x}_{it}$, around average forecasts $\hat{x}_t$ across sessions for each experiment. If there is a pattern, it is for inflation diversity to fall, at least in early sessions of an experiment. Figures C14–17 display time series of $\max_i \hat{x}_{it} - \min_i \hat{x}_{it}$ for each experiment. Vertical lines denote inaugurations of new sessions. Generally, diversity of forecasts is highest at the beginning of an experiment, and there is some tendency for increased dispersion at the inauguration of a new session within an experiment. Only occasionally is there a within-session increase in dispersion.

9. CONCLUDING DISCUSSION

Before our experiments, we were skeptical that chanting "just do it" would solve the time-consistency problem posed by an expectational Phillips curve. Our experiments have softened but not fully arrested our skepticism. A supermajority of experimental sample paths show the monetary authority gradually reaching for the Ramsey value. This might reflect the "just do it" spirit. We think it probably reflects a Phelps–Cho–Matsui monetary authority that imputes an "induction hypothesis" (that is, adaptive expectations) to the private forecasters and that sets out to manipulate private forecasts by its actions. However, there is a big gap between estimated feedback rules and those that would have been chosen by the optimal Phelps planner, who knows the value of citizens' adaptive expectations coefficient. Our policymakers exploit the "induction hypothesis" too slowly, when they exploit it at

all. There are more than enough deviations from Ramsey for us not to take the solution of the time-consistency problem for granted. In addition to occasional backsliding, our experimental economy can be stuck with an incompetent policymaker.

APPENDIX A: INSTRUCTIONS FOR POLICYMAKERS

Today you will participate in an experiment in economic decisionmaking. Various research foundations have provided funds for the conduct of this research. The instructions are simple, and if you follow them carefully and make good decisions, you can earn up to $20, which will be paid to you in cash at the end of this experiment.

You will be assigned the role of a policymaker. In each period of the experimental economy, your job will be to choose the target inflation rate. As a policymaker, you are concerned about the values of inflation and unemployment. However, you can directly affect only the inflation rate.

You will play a series of experimental sessions. An experimental session will consist of a number of experimental periods. At the beginning of each period of an experimental session, you will be asked to choose the target inflation rate. The actual inflation rate will then be determined by adding a stochastic shock to the target inflation rate. This reflects the fact that you, as the policymaker, do not have complete control over the inflation rate.

The stochastic shock is normally distributed and has the mean value equal to 0 and the standard deviation equal to 0.3. This means that approximately 68 percent of the values of the shock will be between –0.3 and 0.3. In addition, approximately 95 percent of the values will be between –0.6 and 0.6. Almost all the values, 99.7 percent, will be between –0.9 and 0.9

At the beginning of each time period, private agents will forecast the inflation rate for that time period. At the end of each experimental period, you will see the average forecasted inflation rate (averaged over the forecasts of all private agents) on your computer screen. You will also see the actual rate of inflation and the rate of unemployment for that experimental time period.

The actual inflation rate and the average forecasted inflation rate (averaged over the expectations of private agents) play a role in determining the rate of unemployment in the economy. The rate of unemployment is calculated in the following way:

$$\text{unemployment} = u^* - (\text{inflation} - \text{average forecasted inflation}) + \text{shock},$$

where u^* is the *natural rate of unemployment*, which prevails in the economy if the actual rate of inflation is equal to average forecasted inflation rate; *average expected inflation* is the rate computed as the average of private agents' expected rates;

and shock is a stochastic *shock* normally distributed, with mean value 0 and the standard deviation equal to 0.3.

At the end of every experimental period, you will also see the payoff that you earned in that period. The payoff is calculated in the following way:

$$\text{payoff} = -0.5\,(\text{inflation}^2 + \text{unemployment}^2).$$

Thus, your payoff decreases with increases in both the inflation and unemployment rates.

At any given experimental period, the probability that the current session will continue for one more period is equal to 0.98. Whether the session is played for one more period is determined in the following way: A random number between 0 and 1 is drawn from a uniform distribution. If the number is less than or equal to 0.98, the current session continues into the next period. If the number is greater than 0.98, the session is over. This number will appear in the last column of your screen at the end of each experimental time period. Once the randomly drawn number is greater than 0.98, the session will automatically be terminated.

You will start every experimental session by running a computer program. The experimenter will give you the name of the program. Once you start the program, you will be prompted to enter the session number. You will enter these numbers in the consecutive order, starting with 1 for the first session, 2 for the second, etc. After entering the session number, you will be prompted to enter the probability that a particular session will end at any given experimental time period. Enter the number 0.98 for this question. Once you have answered these two questions, an experimental session will begin.

Earnings

The experiment will last two hours. If you complete this two-hour experiment, you are guaranteed to receive a $10 payment. Moreover, you can earn an additional $10, for a total of $20.

At the end of each session, the probability of winning a prize of an additional $10 will be computed in the following way.

1. For every time period of the session, the number of period points is calculated by adding 100 points to the payoff you obtained in that time period.
2. The number of total points is calculated by adding the period points earned in all time periods of a given experimental session. If this number is less than 0, it is set equal to 0.
3. The number of maximum points is calculated by multiplying the total number of periods of the session by 100. This number is the number of total points that you would earn in an experimental session if your payoff were equal to 0 in every experimental period.

4. The probability of winning the prize is then calculated in the following way:
   1 – (*maxpoints – total points*) / *maxpoints.*

   Table A1 presents an example of how the total points, maxpoints, and the probability are calculated in a hypothetical experimental session. The length of the session is five experimental periods.

Table A1

Period	Payoff	Period points
1	–20.25	79.75
2	–115.25	–15.25
3	–5.16	94.84
4	–10.37	89.63
5	–30.25	69.75
Total points		318.72

   maxpoints = 100 x 5 periods = 500

   Thus, the probability of winning the prize in this session is

   1 – (500 – 318.72)\500 = 0.64.

   Higher values of your payoff in each time period (lower in absolute terms) result in higher period and total points. Higher values of total points result, in turn, in higher probability of winning the prize.

5. If your total points happen to be less than zero, then your probability of winning the prize in that session is set equal to zero.

At the end of the experiment, one of the sessions that you played will be randomly selected. Each session will have an equal chance of being selected. The session will be selected by running the program draw.exe at the DOS prompt.

Once you type draw and press enter, you will be asked to enter your ID number. Your ID number as the policymaker is 5. Once you have entered it, you will be prompted to enter the total number of sessions played in the experiment. When you enter this number, the computer will randomly choose a number between 1 and the number of sessions played. This number will appear on your computer screen and will indicate the number of the selected session.

The second number that will appear will be a number between 0 and 1, *rand*, drawn from the uniform distribution. You will take that number and compare it to the probability of winning the prize for the selected session. If *rand* is less than or