Central Bank Reputation, Cheap Talk and Transparency as Substitutes for Commitment: Experimental Evidence

John Duffy* and Frank Heinemann[!]

Preliminary and Incomplete
19 August 2014

Abstract: We implement a repeated version of the Barro-Gordon monetary policy game in the laboratory. Our first aim is to assess whether reputation can serve as a substitute for commitment when central banks are free to operate in a discretionary manner. Specifically we ask whether in a repeated game setting, reputational considerations enable the central bank to implement the forward-looking, efficient Ramsey equilibrium, and avoid the discretionary but time consistent one-shot Nash equilibrium involving higher inflation but no difference in output. We find that reputation is a poor substitute for commitment. We then explore in the repeated discretionary environment whether cheap talk, policy transparency or economic transparency by central bankers can yield improvements in the direction of the Ramsey equilibrium relative to the absence of such mechanisms. Of the three mechanisms we consider, cheap talk fares the best in terms of welfare performance relative to the pure discretionary environment. However, the welfare gains from cheap talk decrease over time. Surprisingly, the real effects of supply shocks are better mitigated by a commitment regime in which monetary policy does not respond to shocks than in any discretionary regime. Thus, there is no trade-off between flexibility and credibility. Our findings suggest that discretionary monetary policy is welfare reducing on both dimensions.

Finally, we also consider a pure commitment regime and find that when central bankers are able to pre-commit to a policy, they have no difficulty implementing the equilibrium policy with low inflation.

Our findings suggest that discretionary monetary policy is welfare reducing on both dimensions. there are real welfare-reducing consequences to discretionary regimes, and that various mechanisms intended to improve welfare in this setting have only small effects. Alternatively put, the equilibrium of the one-shot Barro Gordon model provides a reasonable characterization of behavior in the repeated game setting.

*Department of Economics, University of Pittsburgh, Pittsburgh, PA 15260 USA, email: jduffy@pitt.edu
[!] Department of Economics, Technische Universität Berlin, DE-10623 Berlin, Germany, email: frank.heinemann@tu-berlin.de

1. Introduction

Should central bankers commit to a consistent monetary policy or should they be afforded discretion to alter monetary policy depending on current circumstances? This question, as first posed by Kydland and Prescott (1977) and elaborated upon by Barro and Gordon (1983ab) continues to be debated in discussions of monetary policy design. On the one hand, the ability to use monetary policy to respond to various economic shocks as they arise is the main argument in favor of discretionary policy (e.g., Blinder (1999)). On the other hand, it is well known that the ability to commit to a consistent policy course of action can yield welfare improvements over a pure discretionary policy regime through the effect that the commitment policy has on private sector expectations (e.g. Taylor 1999). The latter argument hinges on the reputation that central bankers can achieve from consistently applying a low inflation monetary policy and thus envisions a repeated game setting between the policy maker and the private sector.

In this paper we implement a version of this repeated policy game due to Barro and Gordon in the laboratory with paid human subjects playing the role of the central banker and the private sector. We have several aims in mind. First, we wish to explore whether reputational considerations can serve as a substitute for commitment in a repeated game setting where central bankers lack a commitment device and are free to alter monetary policy each period conditioning on realizations of economic shocks. We compare that regime with one where central bankers have the ability to pre-commit to a course for monetary policy in advance of the formation of private sector expectations.[1] Our experimental findings from these two regimes reveal that reputation is indeed a poor substitute for commitment in that inflation is higher and welfare is lower in the discretionary environment as compared with the commitment regime. Given that finding, our second goal is to evaluate various mechanisms that have been suggested to be useful for overcoming the temptation to succumb to high inflation time-consistent policy making in the repeated game discretionary environment. In particular we explore the role of cheap talk, policy transparency, both cheap talk and policy transparency,

[1] One advantage of conducting a laboratory experiment is that we are able to implement a commitment regime as a theoretical benchmark, while in more natural settings such commitment devices may have credibility problems.

and finally economic transparency in overcoming the inflationary bias under discretion. Many of these mechanisms have been studied in the context of Barro-Gordon type monetary policy game set-ups (see, Geerats (2002) for a survey) and thus the environment we study is appropriate for an analysis of the effectiveness of such mechanisms. We find that of the various mechanisms we study, cheap talk alone results in some welfare improvement relative to the baseline discretionary environment, but the benefits of such a mechanism appear to decline with experience. In the end, we conclude that none of the various mechanisms we examine comes close to achieving the welfare levels of the commitment regime. These findings suggest that there are real welfare-reducing consequences to discretionary monetary policy.

The rest of the paper is organized as follows. Section two reviews the related literature. Section 3 outlines the one-shot and repeated Barro-Gordon model that we implement in the laboratory. Section 4 describes our experimental design and hypotheses and section 5 reports on the main findings from our experiment. Finally, section 6 concludes.

2. Related Literature

Two prior experiments have been conducted that are related to this paper. Van Huyck et al. (1995, 2001) implement a "peasant-dictator" game in the laboratory. In this two-player game, players are randomly divided up between peasants and dictators. Players are randomly paired with one peasant and one dictator in each pairing; each pair then plays an indefinitely repeated game that continues from one round to the next with probability 5/6. Peasants must decide how much of their endowments of beans to eat or plant (invest) yielding new beans in the next period (if there is a next period). The dictator taxes production and can either pre-commit to a tax rate in advance of the peasant's investment (commitment regime) or decide on a tax ex-post, after investments have been made but prior to their realization (discretionary regime). The authors vary the endowments of the peasants and the interest rate earned on investments.[2] They report that

[2] They report that while they tried a treatment where the dictator made cheap talk announcements of intended tax rates (as we do here) they were dissatisfied with the results of this treatment and dropped it from their analysis.

reputation is an imperfect substitute for commitment and that efficiency under discretion is positively associated with the interest rate earned on investment.

Arifovic and Sargent (2003) implement an experimental version of the Kydland-Prescott model using a design that is similar to our own. In their study subjects were divided up between policymakers and private sector forecasters with one policymaker and 3-5 private sector forecasters in each indefinitely lived match. Private sector forecasters move first: their objective was to correctly forecast next period's inflation. These inflation expectations then entered into a Phillips curve relation that determines the extent to which unemployment departs from its natural rate. The central bank moved second. It had noisy control over the actual inflation rate and sought to minimize its expected loss from the equal weighted sum of the square of the unemployment and inflation rates. Arifovic and Sargent study only a discretionary regime and their treatment variables primarily consist of changes in the variance of shocks to the Phillips curve and inflation setting policy rule. By contrast with Van Huyck et al. they report that subjects do learn to coordinate on the first best Ramsey equilibrium consistent with a commitment regime, despite operating in a pure discretionary environment – that is they find that reputation does work as a substitute for commitment. In particular they report that in three fourths of their sessions, policymakers eventually learn the Ramsey zero-inflation policy and stay with that policy for some time, though in several of the economies they report some "backsliding" toward the less efficient Nash equilibrium associated with one-shot pure discretionary regime after the Ramsey equilibrium had been achieved and sustained for some time.

Our experiment complements and builds upon these earlier studies but also differs from them in several respects. First, the model we implement in the laboratory, the Barro-Gordon (1983ab) model differs from the tax policy focus of the Van Huyck et al. study and differs in certain timing respects from the Kydland and Prescott model of monetary policy studied by Arifovic and Sargent. For example, in the model we study, the central banker always learns inflationary expectations *in advance* of setting policy while in the Arifovic-Sargent study this knowledge is only known to policymakers ex-post. Also differently from Arifovic and Sargent, both policymakers and forecasters are fully informed of the model economy, and subjects can play both roles over the course of an experimental session. Similar to Van Huyck et al. we study the case of both commitment and discretion. However we go beyond these two policy regimes and examine choices under the discretionary regime when 1) the policymaker can engage in pre-play cheap talk about intended policy choices, 2) policy choices are made transparent to the private sector at the end of each period, 3) there is both pre-game cheap talk and ex-post policy transparency and 4) there is a regime of economic transparency where the private sector is informed about an economically relevant supply shock prior to forming their expectations (the policymaker is always informed of this shock in advance of setting policy). Thus our experiment goes beyond a comparison of repeated discretionary decision-making versus commitment and begins the important work of evaluating a number of non-reputation-based mechanisms by which it is thought that central bankers might overcome the inflation bias in the repeated, discretionary environments in which they operate.

Our paper is also related to experimental work involving the Stackelberg leadership model as our game always involves sequential moves though our treatments vary whether the first mover is the central bank or the private sector. Experimental research suggests that a sequential move, Stackleberg leadership structure yields higher welfare than does the comparable, simultaneous move (Cournot) model –see, e.g., Offerman et al. (2001). However, that study and other sequential move experimental games are not typically studied in indefinite horizons environments as in this paper.

Cheap talk has also been explored by experimentalists as a mechanism for solving equilibrium coordination problems (see Crawford 1998 for a survey). Duffy and Feltovich (2006) report on experiments where the truthfulness of prior cheap talk messages (the extent of lying) can be evaluated by the receivers of those messages as in our policy transparency treatment. However, once again, that study and most other studies of pre-play communication have not been conducted in indefinitely repeated games where communication might help solve coordination problems arising from folk-theorem results. An exception is Camera et al. 2012 who study indefinitely repeated prisoner's dilemma games where free-form or structured pre-play communication is allowed at various intervals in the supergame. Similar to our findings, Camera et al. report that such communication does not do very much to help subjects achieve the most efficient equilibrium possible as subjects use communication for both benevolent and deceptive purposes.

3. The Model

The model economy we implement in the laboratory is a version of that used by Barro and Gordon (1983ab). We begin with the static version before moving to the repeated (dynamic) version. Within the static environment we consider first the case of discretion and then the case of commitment.

3.1 Static model

Aggregate output, y, is determined according to a Lucas-style aggregate supply function

$$y = y_n + c(\pi - \pi^e) + u,$$

where y_n denotes the natural rate level of output, π denotes the time t inflation rate, π^e denotes private sector expectations of the inflation rate at time t, c is a constant and u is a mean zero random supply shock. This supply function can be derived from the expectations-augmented Phillips curve view of the inflation-output tradeoff and implies that output deviates from its natural rate only in response to unanticipated inflation. The central bank's monetary policy consists of its choice of m, denoting the rate of growth of the money supply, which determines the actual inflation rate according to:

$$\pi = m + v,$$

where v is a mean zero policy disturbance term (e.g., due to changes in the velocity of money or an unanticipated demand shock). The model is closed under the assumption that the private sector has rational expectations, so that $\pi_t^e = \pi_t$ and that the central bank seeks to minimize the time t loss function

$$L = b(\pi - \pi^*)^2 + (y - y^*)^2,$$

where π^*denotes the central bank's desired inflation rate and $y^* > y_n$ denotes the central bank's desired output level which is assumed to exceed the natural rate.

In the discretionary regime, the private sector moves first forming their expectations for inflation, π^e. The central bank is informed of these expectations and takes them as given when minimizing L subject to the expressions for y and π. The central bank's reaction function, given the private sector's expectations for inflation, π^e, is given by:

$$m = \frac{b\pi^* + c(y^* - y_n - u) + c^2\pi^e}{b + c^2}. \tag{1}$$

The private sector is assumed to have rational expectations, and we distinguish whether or not the private sector is informed about supply shocks when forming expectations. If the private sector cannot observe supply shocks, we have that $\pi^e = E(\pi) = E(m) = \frac{b\pi^* + c(y^* - y_n) + c^2\pi^e}{b + c^2}$, or that $\pi^e = \pi^* + \frac{c}{b}(y^* - y_n)$, which implies that

$$m^{NE} = \pi^* + \frac{c}{b}(y^* - y_n) - \frac{c}{b + c^2}u = \pi^{NE} - \frac{c}{b + c^2}u.$$

Thus the policy choice in the Nash equilibrium (NE) of the one-shot discretionary environment involves an inflation rate, π^{NE}, that is greater than the desired level, π^*, by the amount $\frac{c}{b}(y^* - y_n)$, which reflects the well-known inflationary bias of discretionary policy. The third and final term, $-\frac{c}{b + c^2}u$, reflects the central bank's incentive for stabilizing employment by adjusting money supply to supply shocks.

If the central bank releases its information about supply shocks to the private sector before the formation of expectations, then $\pi^e = E(\pi|u) = E(m|u)$, which implies $\pi^e = \pi^* + \frac{c}{b}(y^* - y_n - u)$ and

$$m^{ET} = \pi^* + \frac{c}{b}(y^* - y_n - u) = \pi^{NE} - \frac{c}{b}u.$$

We refer to this regime as one of "economic transparency" and in our experiment we will compare welfare under economic transparency with welfare under the pure (no economic transparency) discretionary and commitment regimes. As is well known from Geraats (2002), a discretionary regime with economic transparency combines a lack of credibility with a lack of flexibility since any policy responses of monetary policy to supply shocks are perfectly foreseen and, therefore cannot affect employment. The inflation bias in the discretionary regime with economic transparency case is the same as in the discretionary regime without economic transparency, but the monetary policy response to supply shocks, is larger when there is economic transparency, causing a higher variation of inflation rates without stabilizing employment. The reason, why money supply responds so strongly to supply shocks is the timing: it responds to private sector expectations that have already responded to the shock. Consequently, the total impact of shocks on inflation gets magnified. This regime combines time-inconsistent levels of inflation with time-inconsistent responses to shocks.

We next consider the commitment regime. In this environment the central bank moves first and pre-commits to set m in advance of private sector expectation formation but may be able to condition this decision on realizations of the shock, u_t. The central bank assumes that the private sector forms rational expectations, $\pi^e = E(\pi|m) = m$. Thus in this setting, the central bank's minimization problem is:

$$\min_m \; E[b(m + v - \pi^*)^2 + (y_n + c(m + v - \pi^e) + u - y^*)^2|u], \text{ s.t. } \pi^e = m.$$

The solution, which we refer to as the one-shot commitment (C) equilibrium is given by $m^C = \pi^*$.

Under commitment, the central bank cannot create surprise inflation. Hence, it cannot stabilize the real economy, and so its best policy is to target the optimal inflation rate irrespective of the supply shock. Comparing this case to the Nash equilibrium under economic transparency, both problems of time inconstancy are solve by the ability to commit: inflation is both stable and at the optimal level. In comparison to the discretionary regime without economic transparency, however, we observe that that while there is no inflationary bias, as $\pi^* < \pi^{NE}$, the inability to stabilize the real sector under commitment causes welfare losses that can be avoided under discretion.[3] This reflects the well-known trade-off between credibility and flexibility of monetary policy. This tradeoff can be mitigated in the repeated (dynamic game) model where the private sector can condition its behavior on the central bank's past responses to supply shocks, so that reputational considerations come into play. We now turn our attention to that setting.

3.2 Dynamic model

Our experiment implements a dynamic repeated-game version of the model described in the previous section. In each period of this dynamic game, the central bank has the option to exploit low inflationary expectations by surprising the private sector with an unexpected high inflation which raises output and thus welfare in that period. However, unexpected high inflation can trigger a rise in future expectations about inflation which is to the central bank's disadvantage. Thus, the central bank has an incentive to keep inflation low, in order to maintain low inflationary expectations. Whether or not long-term reputational considerations for low inflation dominate short-term welfare gains from surprise inflation depends on parameters and on the effect of current inflation on future expectations. The highest incentive to keep inflation low arises if expectations are driven up forever. Here, we focus on conditions under which the efficient outcome, where $\pi_t^e = \pi^*$, can be sustained as an equilibrium of the infinitely repeated game.

In the repeated game, the central bank's objective is to

$$min_{\{m_t\}} E_t \sum_{t=0}^{\infty} \delta^t \, [b(\pi_t - \pi^*)^2 + (y_t - y^*)^2 | u_t]$$

subject to the given processes for π_t and y_t. Depending on the information available to the private sector, e.g., whether they learn ex-post about the policy rule $m(u)$, and provided that the discount factor δ is sufficiently large, the Folk theorem for infinitely repeated games applies so that the set of equilibria ranges from the one-period discretionary Nash equilibrium where $m_t = \pi^{NE} - \frac{c}{b+c^2} u_t$ and private sector

[3] Whether the benefit from a flexible policy response to such shocks outweighs the costs arising from the inflationary bias of discretionary policy depends, of course on the parameterization of the model.

expectations satisfy $\pi_t^e = \pi^{NE}$ to the efficient "Ramsey" solution where the central bank sets $m_t = \pi^* - \frac{c}{b+c^2}u_t$, avoiding the inflation bias but at the same time having the flexibility to stabilize output while private sector expectations satisfy $\pi_t^e = \pi^*$. As this environment involves a multiplicity of equilibria with no clear means of choosing from among this set of equilibria, a laboratory experiment can be informative as to which equilibria agents are likely to coordinate on and under what conditions.

Before turning to the experimental design, we will briefly derive the conditions for the Ramsey solution being equilibrium of the repeated game. We assume throughout that the private sector will learn inflation π_t and employment y_t at the end of the period. Knowing the Phillips curve, it can thereby deduce the supply shock u_t. We need to distinguish, though, whether or not the private sector will get informed about the actual policy m_t and thereby the transmission shock v_t.

If policy is transparent (as assumed by Barro and Gordon (1983b)), expectations in period *t*+1 can be conditioned on the actual relationship between m_t and u_t. The strongest incentive for the central bank to pursue the Ramsey rule is supported by a grim trigger strategy played by private sector agents in which their expectations are initially $\pi_0^e = \pi^*$ in the first period and remain there, as long as the central bank follows the Ramsey rule. If, however, the central bank deviates from this policy in any period τ, the private sector's expectations immediately jump towards the one-period Nash equilibrium $\pi_t^e = \pi^{NE}$ for all $t > \tau$, and the best response for the central bank is to follow the Nash-equilibrium policy m^{NE} in all future periods. Thus, the central bank faces the trade-off between exploiting low expectations and raising employment for one period on the one hand and implementing the maximum equilibrium inflation bias for all future periods as the result. If the central bank deviates in say period 0, it should best respond to $\pi_0^e = \pi^*$ and u_0, which yields $m_0 = \pi^* + \frac{c}{b+c^2}(y^* - y_n - u_0)$ and gives rise to a welfare loss of

$$L_0^D(u_0)$$
$$= E\left[b\left(\frac{c}{b+c^2}(y^* - y_n - u_0) + v_0\right)^2 + \left(y_n + c\left(\frac{c}{b+c^2}(y^* - y_n - u_0) + v_0\right) + u_0 - y^*\right)^2 \middle| u_0\right]$$

$$=b\left(\frac{c}{b+c^2}(y^*-y_n-u_0)\right)^2+b\sigma_v^2+\left(\frac{b}{b+c^2}(u_0-y^*+y_n\right)^2+c^2\sigma_v^2$$ [skip this line in publication]

$$=\frac{b}{b+c^2}(y^*-y_n-u_0)^2+(b+c^2)\sigma_v^2\ .$$

The expected welfare loss associated with a deviation from Ramsey is then given by $L_0^D(u_0)+\sum_{t=1}^{\infty}\delta^t\,E(L^{NE})$, where $E(L^{NE})$ is the prior expected welfare loss in the one-period Nash-equilibrium:

$$E(L^{NE})=E\left[b\left(\frac{c}{b}(y^*-y_n)-\frac{c}{b+c^2}u+v\right)^2+\left(y_n+c\left(-\frac{c}{b+c^2}u+v\right)+u-y^*\right)^2\right]$$

$$=\frac{b+c^2}{b}(y^*-y_n)^2+\frac{b}{b+c^2}\sigma_u^2+(b+c^2)\sigma_v^2\ .$$

This must be compared with the expected welfare loss if the central bank follows the Ramsey rule. In the first period, this loss is given by

$$L^R(u_0)=E\left[b\left(-\frac{c}{b+c^2}u_0+v_0\right)^2+\left(y_n+c\left(-\frac{c}{b+c^2}u_0+v_0\right)+u_0-y^*\right)^2\middle|u_0\right]$$

$$=b\left(\frac{c}{b+c^2}u_0\right)^2+b\sigma_v^2+\left(\frac{b}{b+c^2}u_0-(y^*-y_n)\right)^2+c^2\sigma_v^2$$ [skip this line in publication]

$$=(y^*-y_n)^2+\frac{b}{b+c^2}(u_0^2-2u_0(y^*-y_n))+(b+c^2)\sigma_v^2\ .$$

The period-0 expectation of future losses under Ramsey is

$$E(L^R)=E\left[b\left(-\frac{c}{b+c^2}u+v\right)^2+\left(y_n+c\left(-\frac{c}{b+c^2}u+v\right)+u-y^*\right)^2\right]$$

$$=(y^*-y_n)^2+\frac{b}{b+c^2}\sigma_u^2+(b+c^2)\sigma_v^2\ .$$

Thus, the central bank has *no* incentive to deviate, if and only if $L_0^D(u_0)+\sum_{t=1}^{\infty}\delta^t\,E(L^{NE})\ \geq L_0^R(u_0)+\sum_{t=1}^{\infty}\delta^t\,E(L^R)$, which is equivalent to

$$\frac{b}{b+c^2}(y^*-y_n-u_0)^2-(y^*-y_n)^2-\frac{b}{b+c^2}\left(u_0^2-2u_0(y^*-y_n)\right)\geq\sum_{t=1}^{\infty}\delta^t\left[-\frac{c^2}{b}(y^*-y_n)^2\right]$$

$$\Leftrightarrow\frac{-c^2}{b+c^2}(y^*-y_n)^2\geq\frac{-\delta}{1-\delta}\cdot\frac{c^2}{b}(y^*-y_n)^2\Leftrightarrow(1-\delta)b\leq\delta(b+c^2)\Leftrightarrow\delta\geq\frac{b}{2b+c^2}. \qquad \text{(xx1)}$$

This condition is necessary and sufficient for the Ramsey solution to be an equilibrium under transparent monetary policy.

If policy is intransparent, the private sector cannot perfectly infer whether an increase in inflation is due to the central bank's deviating from the Ramsey rule or to an unfortunate realization of the transmission shock v_t. Here, the parameter restrictions that support the Ramsey equilibrium depend on the distribution of both shocks. In our experiment, we will use uniform distributions with bounded support. This allows us to derive a sufficient condition under which the Ramsey solution is an equilibrium.[4]

Suppose v has a uniform distribution in $[-\mu,\mu]$ and consider the following strategy of forecasters: Expectations start with Ramsey and switch to Nash forever from period $t+1$ onwards, if $\pi_t>\pi^*+\mu$. As long as the central bank plays Ramsey, the probability of expectations switching to Nash is zero. However, the central bank may raise the money supply just enough to exploit the large marginal gains for reducing unemployment from high levels at the risk of a moderate probability of being punished in the future. If μ is large, the probability of detection is small, provided that m exceeds π^* just slightly. The CB can hide behind the shock, which may provide an incentive for deviations from Ramsey. For deriving a sufficient condition that prevents such incentives, first note that the marginal gain from increasing employment in the current period is a concave function of the money supply due to the quadratic loss function. The marginal expected

[4] Henckel et al. (2011) discuss this problem for a normally distributed shock and a welfare function that is linear in output. They assume that expectations switch to the one-period Nash equilibrium for one period if a certain test statistic indicates that the central bank has been cheating with some given probability. However, the test statistic is chosen arbitrarily and it is assumed that the central bank does not strategically game the test statistic. Under these conditions, the Ramsey solution cannot be sustained as equilibrium.

future loss stemming from the probability of being detected, however, is linear due to the uniform distribution of transmission shocks.

If the money supply rises to $m \in (\pi^*, \pi^* + 2\mu)$, the probability of being detected is $prob(\pi > \pi^* + \mu | m) = \frac{m-\pi^*}{2\mu}$.[5] The associated expected welfare loss in the current period, say $\tau = 0$, is

$$E(L|u_0, m) = E[b(m + v - \pi^*)^2 + (y_n + c(m + v - \pi^*) + u_0 - y^*)^2 | u_0]$$

$= b(m - \pi^*)^2 + b\sigma_v^2 + c^2(m - \pi^*)^2 + 2c(m - \pi^*)(u_0 + y_n - y^*) + (u_0 + y_n - y^*)^2 + c^2\sigma_v^2$ [skip]

$$= (b + c^2)(m - \pi^*)^2 - 2c(m - \pi^*)(y^* - y_n - u_0) + (y^* - y_n - u_0)^2 + (b+c^2)\sigma_v^2.$$

Thus, the marginal expected gain from increasing m is $\frac{-\partial \mathrm{E}(L|u_0,m)}{\partial \mathrm{m}} = 2c(y^* - y_n - u_0) - 2(b + c^2)(m - \pi^*)$. The marginal expected loss is $\frac{1}{2\mu}\sum_{t=1}^{\infty} \delta^t \left(E(L^{NE} - L^R)\right) = \frac{\delta}{1-\delta} \cdot \frac{c^2}{2\mu\, b}(y^* - y_n)^2$. A marginal deviation from Ramsey does *not* pay off, if and only if

$$2c(y^* - y_n - u_0) \le \frac{\delta}{1-\delta} \cdot \frac{c^2}{2\mu\, b}(y^* - y_n)^2 \Leftrightarrow 4\mu\, b(y^* - y_n - u_0) \le \frac{\delta}{1-\delta} \cdot c(y^* - y_n)^2.$$

Clearly, the incentive for inflating the economy rises with larger negative supply shocks. Since we assume a bounded support, a sufficient condition preventing deviations from Ramsey is

$$4\mu\, b(y^* - y_n - u_{min}) \le \frac{\delta}{1-\delta} \cdot c(y^* - y_n)^2$$

$$\Leftrightarrow (1 - \delta)4\mu\, b(y^* - y_n - u_{min}) \le \delta c(y^* - y_n)^2$$

$$\Leftrightarrow 4\mu\, b(y^* - y_n - u_{min}) \le \delta[c(y^* - y_n)^2 + 4\mu\, b(y^* - y_n - u_{min})]$$ [skip the last 2 lines in publication]

[5] If money supply is larger, the deviation from Ramsey will be detected for sure. Condition (xx1) ensures that this is not in the interest of the central bank.

$$\Leftrightarrow \delta \geq \frac{4\mu\, b(y^* - y_n - u_{min})}{c(y^* - y_n)^2 + 4\mu\, b(y^* - y_n - u_{min})}, \quad] \qquad \text{(xx2)}$$

where u_{min} is the smallest possible realization of the supply shock. In the experiment, we will make sure that Conditions (xx1) and (xx2) hold.

4. Experimental Design and Hypotheses

Our experimental design consists of six different treatments that vary in the timing of moves and in the information available to participants. However, across all treatments, a number of factors were held constant and we begin with this basic structure.

4.1 Baseline Design

Each session of a given treatment involved 20 subjects with no prior experience with this experiment. The experiment was conducted over networked computers and was programmed using the z-Tree software (Fischbacher 2007). At the start of each session, subjects were randomly divided up into two matching groups of size 10 and there was no further interaction between matching groups so that each session yielded two independent observations. A session for each matching group consists of a number of indefinitely repeated games known as "sequences" with each sequence consisting of an unknown number of rounds.

At the start of each new sequence, subjects in each matching group were randomly divided up into two groups of size 5 and the composition of the group remained constant for all rounds of the sequence. Prior to play of the first round of each sequence, one member of each group was randomly selected to play the role of the central banker, known as the "type A" player, while the other four members of each group were assigned to play the role of the private sector, known as "type B" players.[6] The players' roles remained constant for all rounds of the sequence. At the start of each new sequence, groups were randomly formed

[6] Our choice of having a single central banker and a larger "private sector" of four players follows the setup of Arifovic and Sargent (2003) and reflects that fact that the private sector is considerably larger than the government sector.

anew and the type A player was again randomly chosen from among the membership of the new group, so that there is turnover of central bankers in our environment.

To avoid triggering any pre-conceived notions of the proper role or choices to be made by each player type, we used neutral language and a neutral framing of the model. For instance we did not want players to imagine they were central bankers choosing an inflation rate because they might have used their own experience from the real economy to coordinate on a particular rate, and we wanted the incentive structure of the model to determine their decisions instead. Another change we made in comparison to the model was to cast the loss function in terms of an inflation-unemployment tradeoff so that y_n now refers to the natural rate of unemployment.

[FH: I guess, we should better introduce the model in terms of unemployment already. Relabeling variables is not reader-friendly, is it? JD: Yes, but it seems like a lot of work at this stage. FH: I will do it.]

We told subjects to imagine that the two variables, y_t and $\pi_{t,}$ stand for two "containers" holding varying amounts of water[7]. Subjects were instructed that at the start of each round t =1,2,…, container 1 (unemployment) held w_t "gallons" ("liters") of water where w_t was publicly known to be an i.i.d. random draw each period from a uniform distribution over the interval [120, 160]. The expected value, $E[w_t] = 140$, corresponds to the natural rate of unemployment, y_n, in the model while the supply shock $u_t = w_t - E[w_t]$. Thus, in our parameterization one can think of the supply shock, $u_{t,}$ is an i.i.d. random draw from a uniform distribution with support [-20, 20] and thus having mean 0 and standard deviation $20/\sqrt{3}$ =11.55. The initial amount of water in in container 1 (unemployment) thus consists of both the natural rate level of unemployment and the supply shock, i.e., $w_t = y_n + u_t$. The timing of when or whether players learned the value of w_t is an important element of our treatments. Subjects were further instructed that container 2 (inflation) was initially empty.

[7] The idea of framing a monetary policy game in terms of targeting amounts of water or chips in a container has been used first by Engle-Warnick and Turdaliev (2010).

In our baseline, discretionary policy treatment, the timing of moves was as follows. The four type B players in each economy moved first each submitting a forecast, $\pi_{i,t}^e$, as to how many gallons (liters) of water would be in container 2 at the end of round *t*. They did so without knowing the realized value of w_t, though they *did* know that w_t was an i.i.d. random draw from a uniform distribution on the interval [120, 160] and they were told that $\mathrm{E}[w_t] = 140$. They were also informed about the player A's objective function (as described below), so they knew the player A's (central bank's) target values for inflation, π^*, and unemployment, y^*. After all four player Bs had made their forecasts, the computer program calculated the mean forecast $\pi_t^e = \frac{1}{4}\sum_{i=1}^{4}\pi_{i,t}^e$ for the economy/group and revealed this value to the group's player A – this forecast corresponds to π_t^e in the model. Subjects were instructed that this average forecast value would be added to the amount of water that was already in container 1, so that the amount of water in container 1 now increased to $w_t + \pi_t^e$. Then, the player A alone in each group learned the value of both w_t and π_t^e and the sum $w_t + \pi_t^e$, representing the new total amount of water in container 1. Player A was then instructed to "move" some amount $m \in [0,80]$ of water from container 1 to container 2. This choice corresponds to the policy choice of m_t for period t. In the baseline discretionary treatment, Player Bs do not observe Player A' choice for m_t but it is public knowledge that $m \in [0,80]$. In addition, it was public knowledge to both player types that there was a random, uncontrolled flow of water, v_t, from container 1 to container 2, corresponding to the policy transmission shock of the theory. The value of v_t was publicly known to be an independent random draw each period from a uniform distribution having support [0, 40] and all subjects were instructed that $E[v_t] = 20$. Note that differently from the theory, the transmission shock is not mean zero; this choice was made because the policy action space was $m \in [0,80]$, and we did not want to have inflation $\pi_t = m + v_t$ be negative. Player As do not observe the value of v_t until after they have chosen m_t.

At the end of each period t, the final amount of water in container 1 is thus given by $w_t + \pi_t^e - m_t - v_t$, which correlates to the Phillips curve relationship, in which surprise inflation reduces unemployment. Setting the parameter of the Phillips curve $c = -1$, as we do in all treatments of our experiment, we have that the final amount of water in container 1 corresponds to $y_t = y_n + \pi_t^e - \pi_t + u_t$ while the final amount of water in container 2 corresponds to $\pi_t = m_t + v_t$.

The final amounts of water in containers 1 and 2 were revealed to all subjects in each economy of size 5 at the end of each period as these amounts determined the subjects' payoffs for the round. Specifically, each player type was incentivized to make choices consistent with the objective functions posited by the theory. Type A players were instructed that their point earnings were given by the formula

Player A Points = 6000 – 2 (Final Container 1 amount – 120)^2 – (Final Container 2 amount – 40)^2.

Thus, player A's (central bankers) had as their policy objectives: $y^* = 120$ and $\pi^* = 40$ and the parameter b was set equal to ½. As noted above, the payoff function and the parameter choices for player As were public knowledge to all participants as revealed in the written instructions. Type B players were instructed that their point earnings were given by:

Player B Points = 4000 – ($\pi_{i,t}^e$ – Final Container 2 amount)^2.

Thus, player B's had the simpler task of just forecasting the value for π_t, the amount of water in container 2 correctly.[8] For this reason, the maximum number of points that player Bs could earn each round was 4,000, while player A's, who had the more difficult decision, could earn a maximum of 6,000 points per round. These equations determining players' points were presented to both player types and for ease of understanding, subjects were given payoff tables showing how their choices would convert into points each round. (See the instructions for these payoff tables and other details about what subjects were told) [to be included in an appendix]. The parameterization of the model is justified in section 4.3.

[8] The levels of inflation and unemployment do not enter the payoff functions of forecasters, only the payoffs of the central bankers. We are assuming that inflation forecast accuracy is all that matters to the private sector; given accurate inflation forecasts they would be able to infer the level of output (unemployment) from knowledge of the model and solve any profit or utility maximizing problems that they faced. [elaborate explanation:o.w. individual incentives arise from the possibility to affect aggregate variables, so that rational players would not state their true expectations.]

At the end of each round, all subjects are informed of the final amounts of water in the two containers corresponding to the realizations of y_t and π_t and their point earnings for the round as determined by the expressions given above. In addition, type B players learn the realization of $w_t = y_n + u_t$. As they were told at the outset that the mean value for $w_t = 140 = y_n$, in effect they learn the realization of the supply shock u_t at the end of the period. Whether type B players learn the central banker's choice for m_t depends on the treatment as detailed below

Subjects were instructed that at the end of each round that a six-sided die would be rolled to determine whether the sequence continued with another round; they were further instructed that if a six was rolled, the sequence ended but otherwise the sequence would continue with another round. This constant random continuation probability implements both discounting with factor $\delta = \frac{5}{6}$ and the stationarity associated with an infinite horizon. Subjects were instructed that, depending on the time remaining in the session, a new indefinite sequence might begin.

Subjects were informed that at the end of the session (which could last up to three hours), two sequences would be randomly chosen from among all those played and subjects point earnings from the chosen sequences would be converted into cash at a known and fixed rate. In addition, at the start of each sequence subjects were endowed with 5,000 points to avoid negative payoffs; since two sequences were randomly drawn at the end, the endowment served as subjects' show-up payment.

4.2 Treatments

Our experiment consists of six different treatments that are intended to explore the role (if any) played by reputation, cheap talk and transparency (both policy and economic) on welfare in the repeated discretionary environment relative to an environment where central banks can pre-commit to monetary policy. The six treatments are:

1. **Discretionary policy**: The timing and moves for this baseline treatment are as described in Section 4.1. The private sector, type Bs, move first and are not informed of the realization of the supply shock u_t when forming their expecations, $\pi^e_{i,t}$. The policy maker observes u_t and π^e_t and then chooses m_t. Type Bs never learn the value of m_t *or* $v_{t,}$ but they do learn the final amounts of water in each container, $y_t$and π_t and are informed of the values for u_t and π^e_t at the end of the period.

2. **Commitment:** In this treatment, the central banker, type A moves first, observing the realization of the supply shock, u_t and then choosing m_t prior to the formation of inflationary expectations by the private sector. Thus the commitment environment we study provides the central banker with the ability to respond to supply shocks but also to pre-commit to an inflation policy for the period.

3. **Cheap talk (CT)**: In this treatment, type A players again move first, observing the value of the supply shock u_t and then choose a message of the form: "The amount of water I intend to move from Container 1 to Container 2 is __." In the blank space they enter the value of $m \in [0,80]$ that they want to signal to forecasters in that round. This message is sent to all player Bs in their group. Then the player Bs form their forecasts of inflation for the period, $\pi^e_{i,t}$. Recall that player Bs understand that the final amount of water in container 1, $\pi_t = m_t + v_t$, so the announcement concerning the intended choice for m_t can play a role in coordinating private sector inflationary expectations. To ensure that the message is understood to be cheap talk, subjects are further instructed that "it is up to Player A whether he or she moves as much water as previously announced. Player A can move the announced amount or more or less water."

4. **Policy transparency (PT)**: This treatment has the same timing as the baseline discretionary treatment. The only difference is that the private sector type B players learn the realizations of both m_t and v_t at the end of each round, thus making it transparent as to whether inflation was high (low) due to the transmission shock or due to the policymaker's choice. A transparent policy environment such as in this and the next treatment (cheap talk plus policy transparency) makes it easier to sustain the Ramsey solution as equilibrium of the repeated game.

5. **Cheap talk and policy transparency (CT+PT)**: This treatment combines the cheap talk phase of the cheap talk treatment with the information revealed about monetary policy (m_t and v_t) at the end of each round as in the policy transparency treatment. This treatment thus allows the private sector to evaluate the truthfulness of the central bank's cheap talk announcements providing a potentially more credible means by which the central bank can attempt to manage private sector expectations as compared with cheap talk or policy transparency by themselves.

6. **Economic transparency:** In this treatment, the private sector (player Bs) learn the value of the supply shock u_t at the same time that the central banker (player A) learns it, and prior to forming expectations of inflation for the period. The timing and information is otherwise identical to that of the discretionary treatment.

4.3 Calibration and interpretation

Parameter choices had to satisfy different requirements: subjects must be able to comprehend the relation between different variables, the different equilibria should be sufficiently differentiated, allowing us to detect treatment effects and potentially reject hypotheses, the parameters should allow for an interpretation as a monetary policy game, the continuation probability should allow for several sequences per session and also satisfy the conditions for the Ramsey solution to be an equilibrium. While the levels of the unemployment and inflation target y^* and π^* are mere normalizations, the other parameters affect the interpretation.

We suggest that a period in the experiment corresponds to 2 years, our guess regarding the length of time it takes for monetary policy to have an impact on inflation and unemployment. Thus our continuation probability choice of $\delta = 5/6$ means that the expected duration of a supergame (or a CB policy regime) is 12 years and implies an annual discount rate of about 8 percent. We assume that the difference between natural and efficient unemployment is 2 percent. As we set $y_n = 140$ and $y^* = 120$, the difference between efficient and natural unemployment is 20, so we may think of units of y as equaling

0.1 percent unemployment.[9] We suppose that a 1 percent increase in inflation leads to a .5 percent reduction in unemployment. Setting $c = 1$, thus, implies that units of π must correspond to .2 percent inflation. Since the standard deviation of supply shocks u is 11.55 and a unit of y is .1 percent unemployment, the standard deviation of supply shocks in the experiment is about 1.15 percent unemployment. The transmission shock, v, also has a standard deviation of 11.55 and since a unit of π is .2 percent inflation; the standard deviation of policy transmission shocks in the experiment is 2.3 percent inflation. The parameter choices also imply that the inflation bias in the one-period Nash equilibrium is $\frac{c}{b}(y^* - y_n) = 40$ and, since units of π correspond to .2 percent inflation, corresponds to 8 percent inflation. Finally, we note that the optimal coefficient attached to supply shocks in the one-period Nash equilibrium and in the Ramsey solution is $\frac{c}{b+c^2} = \frac{2}{3}$. Suppose that the economy is hit by a shock that would increase unemployment by 1% without further policy measures ($u = 10$). In our economy, the optimal response would be an increase of π by 6.67 corresponding to 1.33% inflation. This policy reduces the increase in unemployment to .67%.

Note that the condition for existence of a Ramsey equilibrium in the repeated game under policy transparency requires $\delta \geq \frac{b}{2b+c^2} = 0.25$. With $\mu = 20$ and $u_{min} = -20$, the sufficient condition for the Ramsey solution to be an equilibrium in the regimes without policy transparency is $\delta \geq \frac{4\mu\, b(y^* - y_n - u_{min})}{c(y^* - y_n)^2 + 4\mu\, b(y^* - y_n - u_{min})} = \frac{40*40}{(20)^2 + 40*40} = 0.8$. Naturally, it is more challenging than the condition under policy transparency. Our discount factor of $\delta = 5/6$ satisfies both requirements.

4.4 Experimental Hypotheses

Given our parameterization of the model, theory predictions are summarized in Table 1, which reports point predictions if all equilibria give the same predictions or upper and lower bounds for the range of equilibria. The commitment solution is for the CB to set $m_t = m = 20$ for all t since $E[v_t] =$

[9] Recall that by contrast with the theory section, in the experiment *y* refers to unemployment. Thus, positive *u* shocks increase unemployment and call for a higher monetary policy reaction. Surprise inflation lowers *y*. [footnote can be removed, if we adjust the model]

20, and therefore $E[\pi_t] = m + E[v_t] = 40 = \pi^*$. As laid out in Section 3.1, the CB should not respond to supply shocks, because any such response would be reflected in forecasts and, thus, not affect employment. It follows that inflation varies only with the transmission shock, $\text{Std}[\pi] = \text{Std}[v] = 11.55$, while the variance of unemployment is the sum of the variances of the two shocks, or $\text{Std}[y]=\sqrt{\text{Var}[v] + \text{Var}[u]} = 16.33$. Expectations should be rational and equal π^*, so that the standard deviation of $\pi_t^e - \pi_t$ should simply reflect the standard deviation of the policy transmission shock, $\text{Std}[v] = 11.55$. Expected CB welfare is given by

$$6000 – 2E((140 – 120 - u - v)^2) – E(v^2) = 5200 - 2\text{Var}(u) - 3\text{Var}(v) = 4533.3.$$

Table 1: Experimental Predictions

Predict \ Treat.	Commitment	Discretion	Cheap Talk	Policy Transparency	CT + PT	Economic Transparency
Inflation bias	0	[0,40]	[0,40]	[0,40]	[0,40]	[0,40]
Response of m_t to supply shock, u_t.	0	2/3	2/3	2/3	2/3	2/3 2 *
Response of m_t to exp inflation, π_t^e.	n.a.	2/3	2/3	2/3	2/3	2/3
Response of π_t^e to supply shock, u_t	0	n.a.	0	n.a.	0	2
St.Dev. of π_t	11.55	13.88	13.88	13.88	13.88	25.82
St. Dev. of y_t	16.33	12.17	12.17	12.17	12.17	16.33
St. Dev. of $(\pi_t^e - \pi_t)$	11.55	13.88	13.88	13.88	13.88	11.55
CB Welfare	4533	[3111, 4711]	[3111, 4711]	[3111, 4711]	[3111, 4711]	[2400, 4000]

Note: * Response of *m* to supply shock in ET should be 2/3 if response to expectations is controlled for. Since expectations respond to supply shock as well, the total response of *m* to *u* should be 2.

In the case of discretion, we have multiple equilibria ranging from the one-period Nash with an inflation bias of 40 to the Ramsey solution without inflation bias. In all equilibria, the CB responds to supply shocks and stabilizes employment with a coefficient of $\partial m/\partial u = 2/3$. The same coefficient applies to the central bank's response to inflation expectations, as given by the response function (1). Note however that in a

Ramsey equilibrium forecasts are constant and, thus, there should be no detectable variation to which the central bank responds. As subjects are most likely not exactly in equilibrium, the best response of the central bank with respect to "near Ramsey" expectations might still be zero. In the Result section, we will analyze whether subjects respond in an optimal way to the actual decision of other players. For rational expectation, actual choices of m should range from $\pi^* - E(v) + \frac{2}{3}u_{min} = 6.67$ under Ramsey to $\pi^* + 40 + E(v) + \frac{2}{3}u_{max} = 73.33$. Our subject central bankers can choose values of $m \in [0{,}80]$, which contains all of these vales in the interior, allowing us to test point predictions associated with the most extreme equilibria.

The stabilization of employment raises the standard deviation of inflation to $\text{Std}[\pi] = \sqrt{\text{Var}[v] + \frac{4}{9}\text{Var}[u]} = 13.88$, while it reduces the standard deviation of unemployment to $\text{Std}[y] = \sqrt{\text{Var}[v] + \frac{1}{9}\text{Var}[u]} = 12.17$. Since inflation expectations are constant in any equilibrium, the standard deviation of $\pi_t^e - \pi_t$ should equal the standard deviation of inflation.

In the Ramsey equilibrium, expected CB welfare is

$$6000 – 2E\left(\left(140 – 120 - \frac{1}{3}u - v\right)^2\right) – E\left(\left(v + \frac{2}{3}u\right)^2\right) = 5200 - \frac{6}{9}\text{Var}(u) - 3\text{Var}(v) = 4711.1.$$

In the Nash equilibrium, however, expected welfare is lower, because of the inflation bias and is given by

$$6000 – 2E\left(\left(140 – 120 - \frac{1}{3}u - v\right)^2\right) – E\left(\left(40 + v + \frac{2}{3}u\right)^2\right) = 3600 - \frac{6}{9}\text{Var}(u) - 3\text{Var}(v) = 3111.1.$$

Notice that the discretionary regime admits a range of possible welfare values for the CB that includes the commitment welfare value in its interior.

Policy transparency, cheap talk, and policy transparency plus cheap talk have no impact on the set of equilibria. While cheap talk announcements should be ignored in equilibrium, they might transmit information about the supply shock to the private sector, which justifies testing the equilibrium prediction of expectations being unresponsive to supply shocks and announcements.

Under economic transparency, on the other hand, the private sector learns the supply shock prior to forming expectations. As laid out in Section 3.1, this raises the equilibrium coefficient by which money supply, inflation, and inflation expectations respond to supply shocks to $c/b = 2$ without stabilizing employment, which results in $\mathrm{Std}[\pi] = \sqrt{\mathrm{Var}[v] + 4\mathrm{Var}[u]} = 25.82$, while $\mathrm{Std}[y] = 16.33$ as in the commitment regime. Because the private sector can predict monetary policy responses to supply shocks, the standard deviation of $\pi_t^e - \pi_t$ should simply reflect the policy transmission shock, $\mathrm{Std}[v] = 11.55$. As in the baseline discretionary treatment, there is a set of equilibria, ranging from the one period Nash to a constrained efficient solution, in which there is no inflation bias, but no stabilization of employment either. Welfare in the one-period Nash is given by

$6000 – 2E((140 – 120 - u - v)^2) – E((40 + v + 2u)^2) = 3600 - 6\mathrm{Var}(u) - 3\mathrm{Var}(v) = 2400$. In the most efficient equilibrium, welfare is

$$6000 – 2E((140 – 120 - u - v)^2) – E((v + 2u)^2) = 5200 - 6\mathrm{Var}(u) - 3\mathrm{Var}(v) = 4000.$$

Beside the point predictions from equilibria, we have some more fundamental hypotheses that can be divided up between those pertaining to CB behavior, those pertaining to the behavior of the private sector and aggregate outcomes involving both sets of actors.

Central Bank Hypotheses

Hypothesis 1. Repetition in the discretionary regime serves as a substitute for commitment regarding monetary policy.

By this we simply mean that in the repeated game the CB behaves as if s/he were operating under a commitment regime and (a) produces a similar average money supply. The counter-hypothesis of a higher money supply arises from the inflation bias associated with the less efficient equilibria under discretion.

(b) The second part of this hypothesis states that monetary responses to supply shocks are the same. Here, the counter-hypothesis comes from the positive coefficients in the equilibria of discretionary regimes.

Hypothesis 2. The discretionary regimes with and without cheap talk, policy transparency, or economic transparency produce the same central bank behavior.

More precisely, we test, (a) whether these treatments lead to the same average money supply and (b) to the same central bank responses to to expected inflation and supply shocks.

Private sector hypotheses

Hypothesis 3: Average forecasts of the private sector are unbiased and the distribution of forecast errors reflects the unpredictable fluctuations of inflation only.

This hypothesis follows from the assumption of rational expectations. We will test (a) whether biases are null and, if not, whether they differ across treatments, and (b) whether the fluctuations of expectation errors are the same across treatments if we substract the part that can be explained by transmission shocks and central bank responses to supply shocks unknown to the private sector.

Hypothesis 4: The distribution of individual forecasts around the average is the same across treatments.

We will test how dispersed individual expecations are. While in equilibrium all agents hold the same expectations, in the experiment expectations will differ across subjects. These differences may be related to the treatment, even if the average forecast errors are the same across treatments. Related to this, the payoffs for type-B players might differ across treatments for two reasons: (i) average forecasts may be less precise in some treatments than in others (Std. of $\pi_t^e - \pi_t$) and (ii) individual forecasts may be more or less dispersed (Std. of $\pi_t^i - \pi_t^e$). Theory predicts that B-players should have the highest payoffs in the commitment treatment and under economic transparency, and lower (but equal) payoffs in the other discretionary treatments.

Hypothesis: 5 Under the economic transparency regime alone, inflation forecasts should respond to supply shocks. Cheap-talk announcements should be ignored.

The latter hypothesis follows from the different timing that is assumed in the economic transparency environment. The counter-hypothesis arises from possible information transmission in cheap-talk treatments

Finally, we have some hypotheses regarding *aggregate outcomes*:

Hypothesis 6: Under commitment and economic transparency, output volatility is larger than in the other discretionary regimes.

Here, we test whether output volatility is the same between ET and commitment, whetehr it is the same across the other treatments, and whether there are differnces between these two groups of treatments.

Hypothesis 7: Compared with discretionary treatments , the standard deviaton of inflation should be higher under economic transparency and lower under commitment.

Hypothesis 8: Repetition in the discretionary regime serves as a substitute for commitment regarding (central bank) welfare.

While the inflation bias in discretionary regimes may be higher than under commitment, the flexibility with which central banks can respond to supply shocks under discretion, may reduce employment fluctuations. Thus, there are two opposing effects for the final level of central bank welfare. This can be seen in Table 1, where equilibrium welfare under commitment is strictly in between the lowest and highest welfare levels associated with equilibria in discretionary treatments.

Hypothesis 9: Under economic transparency, central bank welfare is lower than under commitment or in the baseline discretionary treatment.

For economic transparency the welfare levels in all equilibria are lower than under commitment. The only welfare relevant difference between economic transparency and the baseline discretionary treatment is that economic transparency does not allow for stabililizing output (in theory). This justifies our last hypothesis.

Unless stated otherwise, we use non-parametric tests based on average observations from a matching group. For comparing levels. We apply this conservative treatment of data, because behavior of different subjects from the same matching group need not be independent. For point predictions arising from theory, we use the two-sided Wilcoxon matched pairs test, for comparing different treatments, we use the two-sided Mann-Whitney U-test.

Coefficients are tested on the basis of confidence intervals from regressions with fixed effects for matching groups.

5. Experimental Findings

5.1 Money supply

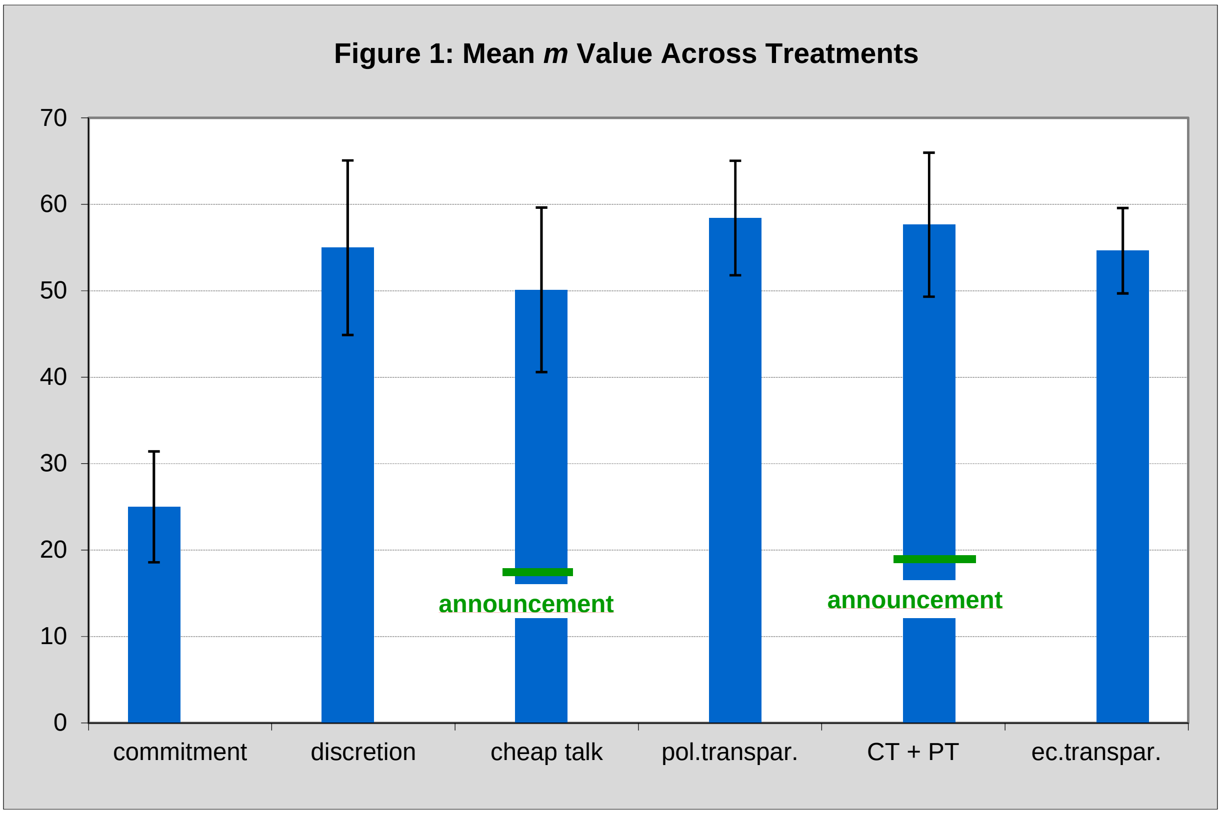

Finding 1: Inconsistent with Hypothesis 1a, reputation does not appear to serve as a substitute for commitment in any of the discretionary regimes.

Evidence in support of Finding 1 is presented in Figure 1 which shows the mean choice of *m* over all supergames of all sessions of each of our six treatments. Also included is a one-standard error bar and the mean announced value of *m* in the two treatments involving pre-game communication. The Figure shows clearly that the mean choice of *m* in the commitment treatment is indistinguishable from the Ramsey solution *m*=20, whereas the mean value of *m* in the other 5 treatments is significantly greater than 20.

2-sided Mann-Whitney U-tests show that money supply under commitment is smaller than in any of the other treatments (p<1%), while there are no significant differences between these other treatments (p>5%). A 2-sided Wilcoxon matched pairs test cannot reject that *m*=20 under commitment (p=15%), but rejects this efficient value for all other treatments (p<1%). The hypothesis that average money supply is equal to the point prediction of the one-period Nash equilibrium (*m*=60) can be rejected for cheap talk (p=3.5%) and economic transparency (p=1.6%), but not for the other discretionary treatments (p>10%). This shows that the one-period Nash equilibrium may be highly relevant in a repeated game, even if all theoretical conditions for the Ramsey solution being an equilibrium hold. Furthermore, neither cheap talk, nor policy transparency, nor economic transparency are effective in reducing the inflation bias, which is consistent with Hypothesis 2a.

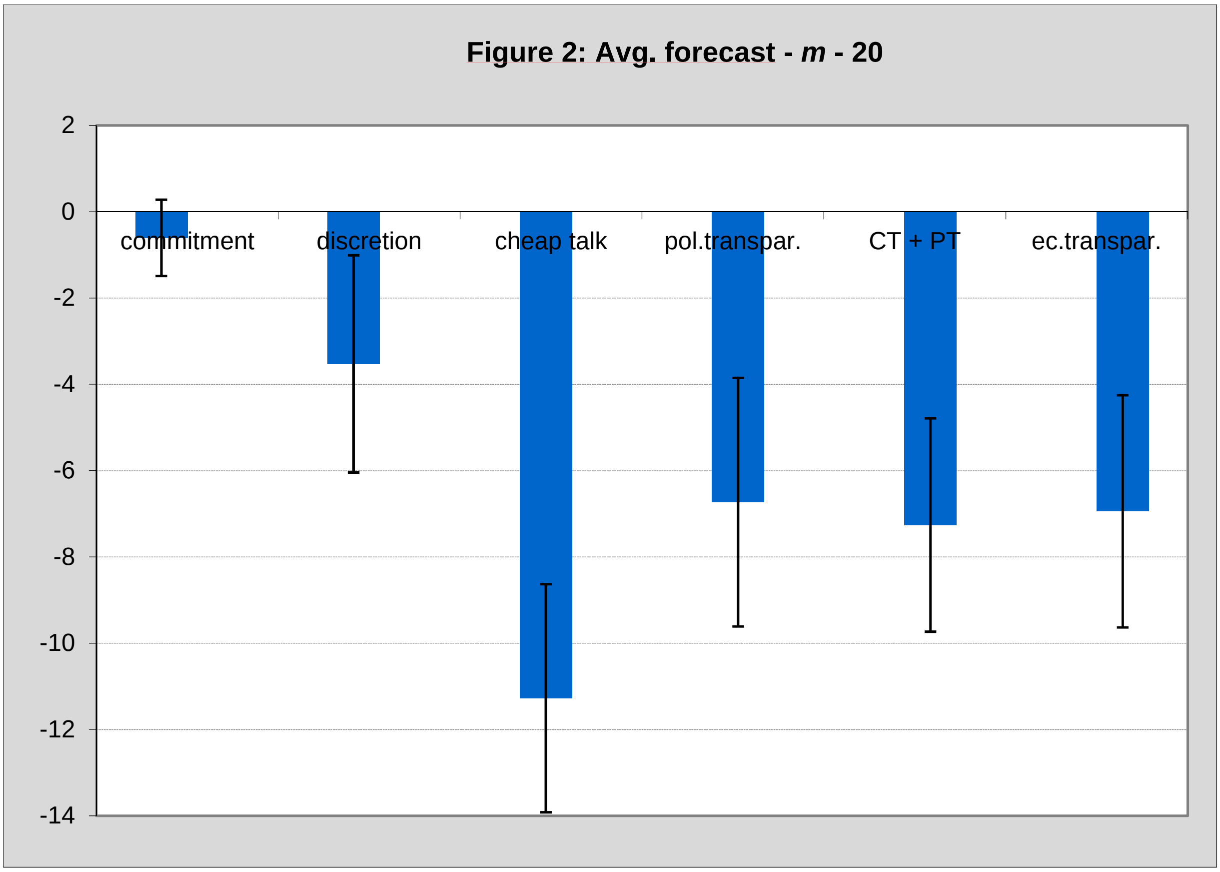

Finding 2: Inconsistent with Hypothesis 3a, private sector expectations are biased in the five discretionary treatments, though not in the commitment treatment.

Evidence in support of Finding 2 is presented in Figure 2. Since inflation equals money supply plus a transmission shock with average value of 20, unbiased forecasts should equal *m*+20. While this hypothesis cannot be rejected for commitment (p=7.8%), we can clearly reject it for all discretionary treatments (p<1%). Comparing the different treatments we find that under commitment, the expectation bias is smaller than in all discretionary treatments (p<1%), under cheap talk, the bias is larger than in the other discretionary treatments (p<1.1%). PT, PT+CT, and ET produce similar biases (p>40%) and they are all larger than under baseline discretion (p<3%). The existence and order of expectation biases comes at a surprise and we will discuss possible reasons below.

As an immediate consequence of biased expectations, average unemployment is smaller than the natural rate. Figure 3 shows the average final Container-1 amount that represents unemployment in our model. While the target rate is 120, the natural rate arising on average in any rational expectations equilibrium is 140. Since private sector expectations fall short of average inflation rates, the unemployment rate deviates from the natural rate towards the central banks' target rate (except for commitment). Under cheap talk the effect is so strong that average unemployment is closer to the target than to the natural rate.

Figure 3: Average unemployment (final Container 1 amount)

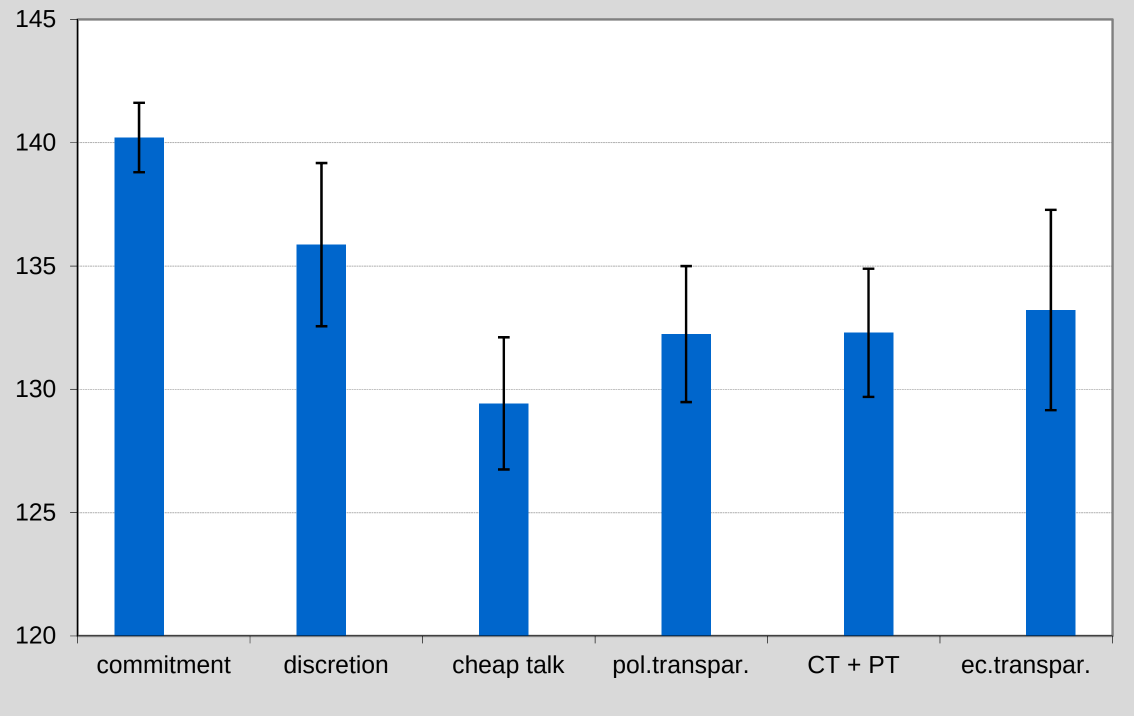

Besides the different inflation biases across treatments and the different abilities of central banks to stabilize employment, the systematic deviations of average unemployment from the natural rate provide a third and unexpected factor influencing central bank welfare. For testing Hypothesis 8, we use two different measures: we compare the actual payoffs of our central bankers as a measure of central bank welfare, but we also compare the payoff with the payoff that the central bank would have achieved in the Ramsey equilibrium under optimal responses to shocks. The advantage of the second measure is that it eliminates fluctuations in payoffs stemming from shocks. The first measure, however, allows a direct test of the point predictions arising from the various equilibria.

Figure 4: actual welfare / welfare in Ramsey equilibrium

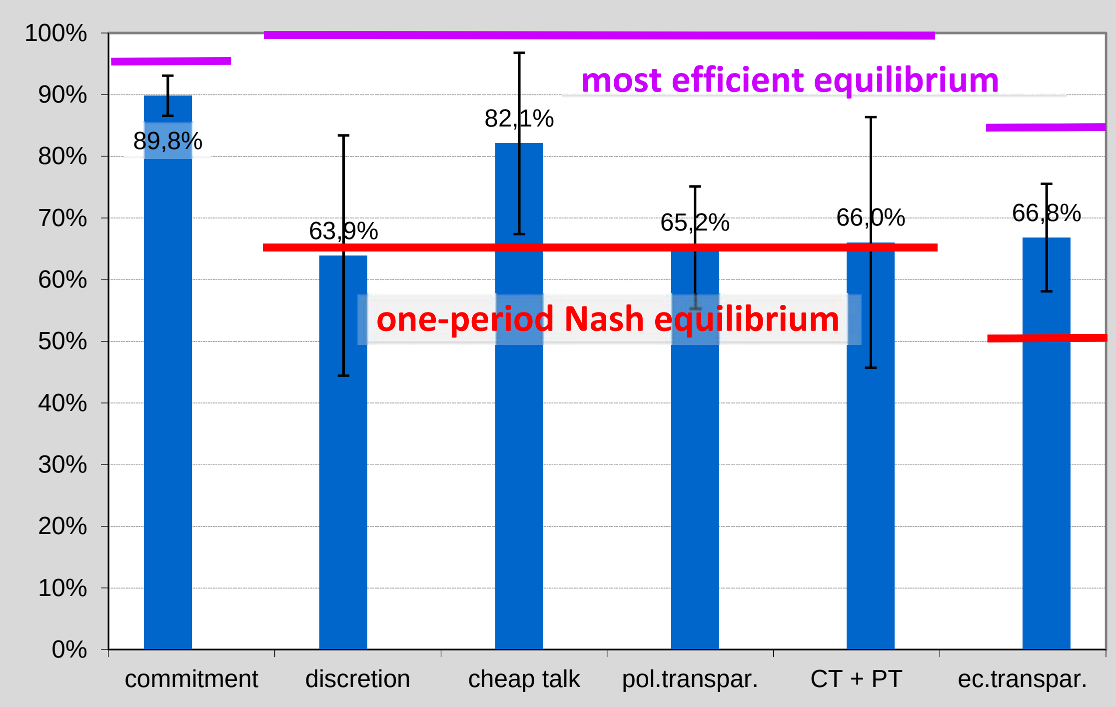

Figure 4 shows that average welfare is below the welfare level associated with optimal policy in all treatments (p<4%). However, since our subject central bankers are not perfect, they also make mistakes under commitment, so that the payoffs of A-players in this treatment is smaller than predicted by equilibrium (p<1%). If we compare the achieved payoffs between different treatments, we cannot reject

Hypothesis 8 for the cheap talk treatment. Here, the achieved payoffs are not smaller than under commitment (p=23%), while they are significantly smaller for the other four discretionary treatments (p<1.1%). Comparing the payoffs for cheap talk with the other discretionary treatments directly, the evidence is mixed: central bank payoffs are higher under cheap talk than for policy transparency (p=2.1%) and economic transparency (p=4.99%), but not significantly different from baseline discretion (p=8.3%) and PT+CT (p=13%). Comparing welfare with the predictions of the one-period Nash equilibrium, we can clearly reject equality under cheap talk (p=1.6%), but not for the other discretionary treatments (p>80%) except economic transparency.

Regarding Hypothesis 9, we can reject that central bank payoffs under economic transparency equal those under commitment (p<1%), but we cannot reject that they are equal to the payoffs in the baseline discretionary treatment (p=100%). For economic transparency, the one-period Nash equilibrium and the most efficient equilibrium are both associated with lower payoffs than the corresponding equilibria under baseline discretion. This is because monetary policy cannot stabilize employment in equilibrium if its responses to supply shocks are anticipated. However, the forecasts of our B-players respond less to supply shocks than money supply, which results in some stabilization of output. In effect the stabilization of output is as good as under baseline discretion, so that we see similar levels of central bank welfare in both treatments.

In order to test hypothesis concerning the players' responses to each other and to supply shocks, we ran treatment-specific regressions with fixed-effects for matching groups. The main regression for central bank behavior was $m_t = \alpha + \beta_1 u_t + \beta_2 avg.f_t$, with the additional restriction β_2=0 for the commitment treatment (where the central bank decided before knowing the private sector's forecast. Regression results are displayed in Table 2

Table 2: Central Bank Behavior

Treat. / Parameter	Commitment	Discretion	Cheap Talk	Policy Transparency	CT + PT	Economic Transparency

α	-35.91*** (6.49)	-2.75 (7.68)	-49.29*** (5.79)	-23.69 (9.43)	-30.35*** (9.97)	-32.21*** (13.34)
β_1	0.433*** (0.047)	0.349*** (0.043)	0.504*** (0.038)	0.419*** (0.039)	0.485*** (0.059)	0.399*** (0.11)
β_2	-	0.137** (0.063)	0.487*** (0.048)	0.210** (0.086)	0.302*** (0.082)	0.456*** (0.084)
R^2	0.18	0.35	0.57	0.22	0.33	0.38

Note: *** significant at 1%, ** significant at 5%, numbers in parantheses are estimated standard deviations.

Coefficient β_1 is positive and highly significant in all treatments (p<1%). This tells us that central bankers are responding to supply shocks, even in the commitment treatment, where this is not optimal provided that forecasters are rational. Whether central bankers and forecasters respond optimally towards each other will be analyzed further below. Comparing the 95% confidence interval with equilibrium predictions, we find that the central bankers' responses to supply shocks are significantly smaller than the equilibrium prediction of 2/3 is all discretionary treatments. Central bankers also respond to average forecasts with coefficients that are positive, but smaller than predicted by the one-period Nash equilibrium (p<5%).

For analyzing the behavior of type B players, we estimate a regression explaining the individual forecasts f_t^i by the current information of the private sector in the particular treatment, lagged own expectations, lagged forecast error, and lagged average forecast. While the lagged forecast error can identify eventual learning behavior, the coefficient on the lagged average forecast provides a measure of convergence. $f_t^i = \alpha + \delta_1 m_t + \delta_2 ann_t + \delta_3 u_t + \beta_1 f_{t-1}^i + \beta_2 (\pi_{t-1} - f_{t-1}^i) + \beta_3 avf_{t-1}$.

Table 3: Private Sector Forecasts

Treat. / Parameter	Commitment	Discretion	Cheap Talk	Policy Transparency	CT + PT	Economic Transparency
α	18.81*** (0.35)	13.97*** (1.88)	36.52*** (3.71)	29.45*** (6.07)	46.29*** (5.47)	-20.46*** (10.78)
δ_1	0.97*** (0.02)	-	-	-	-	-

δ_2	-	-	0.42*** (0.04)	-	0.18*** (0.05)	-
δ_3	-	-	-	-	-	0.52*** (0.07)
β_1	0.07 (0.04)	0.42*** (0.06)	0.16 (0.10)	0.32*** (0.06)	0.13*** (0.06)	0.07 (0.10)
β_2	0.05*** (0.01)	0.29*** (0.01)	0.13** (0.05)	0.14*** (0.03)	0.10*** (0.02)	0.12*** (0.02)
β_3	-0.03 (0.03)	0.40*** (0.05)	0.09 (0.07)	0.27*** (0.05)	0.17 (0.10)	0.17*** (0.09)
R^2	0.80	0.74	0.38	0.58	0.41	0.25

Note: *** significant at 1%, ** significant at 5%, numbers in parantheses are estimated standard deviations.

Forecasters respond to all of their information. In the Commitment Treatment, theory predicts $\delta_1 = 1$, which we cannot reject. Under economic transparency, subjects respond to supply shocks, but the coefficient is significantly smaller than the theoretical prediction of 2. Note that in both cheap-talk treatments, forecasts respond to the announcements made by the central bank. If cheap talk is combined with policy transparency, the coefficient is much smaller than without policy transparency, but they are both significantly positive. This shows that our subject central bankers are successful in affecting expectations by non-binding announcements, in particular when these announcements cannot be fully reviewed by forecasters. This may come as a surprise and it explains why we observe the largest bias between actual and expected inflation in the cheap-talk treatment without policy transparency. We will further analyse the announcements below, and analyse whether central bankers try to exploit their effect on expectations.

In all treatments, subjects respond to the lagged forecast error. The positive coefficient β_2 indicates learning. If we just look at the difference $\beta_1 - \beta_2$, it is is significantly different from zero only in the treatments of discretion and policy transparency. In the same treatments, and under economic transparency, lagged average forecasts have a positive effect.

Next, we look at the stability of (un)employment across treatments. Due to the combined responses of monetary policy, announcements, and expectations (when it applies) to supply shocks, there may be significant differences in the stability of unemployment across treatments. Just looking at standard deviations of employment, we find that it is significantly larger than predicted by theory in all treatments. In fact, our subject central bankers even contribute to fluctuations of employment, because standard deviations are usually higher than they would be for a constant money supply. Table 4 displays the average standard deviation of unemployment across different matching groups within a treatment.

Table 4: Average Standard Deviation of Unemployment

	Commitment	Discretion	Cheap Talk	Policy Transparency	CT + PT	Economic Transparency
$\mathrm{std}(y_t)$	16.56	19.48	17.73	18.66	19.03	20.55

[actually, these numbers measure employment fluctuations within a matching group. If the two groups within a matching group have very different expectation biases (so that their employment **levels** differ), the std(y) is high, even though it may be low within each sequence. Maybe, we should provide measures of std(y) within sequence – but this comes with other problems: (i) how to aggregate these numbers for a matching group, (ii) the variation that is due to changes in policy makers is eliminated – but may be an important argument in favour of rules, see below.]

Two-sided Wilcoxon matched pairs test show that fluctuations of employment are sugnificantly larger than 16.88 (the value for no stabilization) under baseline discretion (p=2.3%), economic transparency (p<1%), and the combination of cheap talk and policy transparency (p=3.9%). In the other treatments, there is no significant difference. The reason for these high variations of unemployment in spite of central bankers responding to supply shocks in the right direction are (i) central bankers are changing over time and different central bankers are pursuing different policies, (ii) forecasters and their forecasts are changing over time and fluctuating expectations contribute to fluctuations in employment, even if central bankers respond to expectations in an optimal way,[10] and (iii) central bankers respond with suboptimally small coefficients to supply shocks and expectations, so that the combination of both may leave unemployment with more fluctuations as under constant expectations and constant money supply.

Comparing the stability of unemployment between treatments, 2-sided Mann-Whitney-U tests show that employment fluctuates less under commitment than under baseline discretion (p=2.1%) and economic transparency (p<1%). Cheap talk also leads to lower fluctuaions than econmic transparency (p=2.1%). The other pairwise comparisons do not yield significant differences (p>10%). It is remarkable that discretionary monetary policy leads to higher fluctuations of employment than commitment (if there is a difference at all). The reason is simply that in the commitment regime inflation expectations and money supply move almost 1 to 1 (see the coefficient δ_1 in Table 3). Hence, the only remaining impact on employment comes from exogenous shocks, while in discretionary treatments the attempts of different central bankers to improve policy lead to higher and unsystematic fluctuations between actual and expected inflation that contributes to destabilizing employment. This result was also surprising us. It provides a strong argument in favor of rule-based monetary policy. The fact that cheap talk leads to the second lowest standard deviation of employment indicates that the responses of the private sector to announcements could have a small stabilizing effect. Of course, this requires that actual policy is correlated with announcements, as is indeed the case (as ww will show below). Howver, since the differences in std(y) between cheap talk and most other treatments are not significant, we cannot draw any conclusions for eventual stabilizing effects of cheap talk. This topic requries further exploration.

To be continued…

[10] Note that expectations are constant in equilibrium (except for economic transparency) but fluctuate quite a lot in the experiment. An increase in expected inflation has the same effect on the economy as an adverse supply shock.

6. Conclusions

To be written

References

Arifovic, J. and T.J Sargent (2003) "Laboratory Experiments with an Expectational Phillips Curve," in D.E. Altig and B.D. Smith (eds.) *Evolution and Procedures in Central Banking*, Cambridge Univ. Press.

Barro, R.J., and D.B. Gordon (1983a) "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy* 91(4), 589-610.

Barro, R.J., and D.B. Gordon (1983b) "Rules, Discretion and reputation in a Model of Monetary Policy," *Journal of Monetary Economics* 12, 101-121.

Blinder (1999)

Engle-Warnick, J., and N. Turdaliev (2010) "An Experimental Test of Taylor-type Rules with Inexperienced Central Bankers," *Experimental Economics* 13, 146-166.

Geerats, P. (2002) "Central Bank Transparency," *Economic Journal* 112, F532-F565.

Henckel, Timo, Gordon D. Menzies, Nick Prokhovnik, and Daniel J. Zizzo (2011) "Barro-Gordon revisited: Reputational equilibria with inferential expectations," *Economics Letters* 112, 144-147.

Kydland, F., and E.C. Prescott (1977) "Rules Rather than Discretion: the Inconsistency of Optimal Plans," *Journal of Political Economy* 85, 473–491.

Taylor (1999)

Van Huyck, J.B., R. Battalio and M.F. Walters (1995) "Commitment versus Discretion in the Peasant-Dictator Game," *Games and Economic Behavior* 10(1), 143-171.

Van Huyck, J.B., R. Battalio and M.F. Walters (2001) "Is Reputation a Substitute for Commitment in the Peasant-Dictator Game?" working paper, Texas A&M University.