Managing expectation in the New Keynesian model

Robert G. King[a], Yang K. Lu[b*]

[a] Boston University;

[b] Hong Kong University of Science and Technology;

February 2019

Abstract

We study the optimal monetary policy in a setting where the private sector is forward-looking and learning about the type of central bank in place. We consider two types of central bank, one patient type that can commit and the other type is myopic and cannot commit. Being able to commit or not, the central bank in place chooses inflation policies optimally, taking into account the learning and rational expectation of the private sector. We show that the equilibrium can be obtained as a solution to a recursive optimization of the committed type in which the actions of the non-committed type are subject to an incentive compatibility constraint. The numerical solution to a calibrated model reveals that the committed central bank with good initial reputation adopts policies similar to the standard solution under full commitment, whereas the committed central bank with poor initial reputation aims at building reputation with anti-inflation policies that involve real costs. If the non-committed central bank with good initial reputation is in place, there will be lengthy real stimulations with gradually rising actual and expected inflation, followed by stagflation when the history of positive inflation surprises depletes the central bank's reputation.

JEL classification: E52, D82, D83

*Corresponding author. Department of Economics, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong SAR. Tel: +852 23587619. E-mail address: yanglu@ust.hk. We would like to thank the audience at various conferences and seminars for helpful comments and discussions. All errors are ours. We acknowledge the financial support from the RGC of HKSAR (GRF HKUST-16504317).

1. Introduction

Since the rational expectation revolution of 1976, it has become conventional for macroeconomists to discuss the historical behavior of monetary policy, inflation and real activity using various combinations of the following six *concepts*:

(1) shifting monetary policy regimes;
(2) optimal monetary policy, with and without commitment;
(3) management of expectations;
(4) private sector learning about the nature of the monetary policy authority;
(5) reputation as a determinant of monetary policy actions;
(6) inflation targets.

Among the historical experience that macroeconomists have sought to understand are six *phenomena*:

(1) lengthy intervals of low and relatively stable inflation and real activity
(2) anti-inflation policies, both gradual and abrupt, involving real costs;
(3) time-varying responses of inflation to the relative price of energy;
(4) lengthy real stimulations with gradually rising actual and expected inflation;
(5) rising inflation accompanied by declining real activity (stagflation);
(6) high inflation with little relationship to real activity.

Our work makes two contributions, building a model environment in which each of the 6 elements are present and illustrating how the model can capture each of the 6 outcomes.

1.1. Modelling

We develop a New Keynesian (NK) model in which there are stochastic shifts in regime, with a random event determining the type of the policy authority, who may able to commit or not. Within a regime, a committed and trustworthy policy authority implements an optimally designed inflation target that takes into account: (i) a conventional objective over inflation and real activity; (ii) the forward-looking nature of inflation in the NK model; and (iii) private sector expectation formation that optimally processes noisy information about the policy authority's type.[1] The system of time-varying inflation targets is calculated using a recursive approach to optimal policy design, so that history dependence is captured using a pseudo state variable (commitment multiplier). The target inflation rate also depends on an exogenous energy price shock and on a reputational state variable that summarizes the private sector's beliefs of the policy authority's type. Learning is incomplete because inflation outcomes are stochastic given the policy authority's action, which is not directly observable although it is identical to the target of a trustworthy committed type. In designing policy, the committed authority understands the evolution of private sector beliefs about its type and the consequences that these have for inflation expectations. Management of expected future inflation thus depends on the likelihood that each type of policy authority will be in place in the future and the actions that a future policy authority will take given the circumstances encountered. The recursive optimal policy design problem for the committed type that we solve is essentially a dynamic principal-agent problem in which the actions of the alternative type – that cannot commit – are subject to an incentive compatibility constraint.

1.2. Applications

Our framework permits us to address important topics in the normative and positive analysis of monetary policy when there are credibility concerns. On the normative side, the solution to our recursive optimal policy problem answers the question: "what is the best policy for a committed authority when there are private sector credibility concerns stemming from the fact that commitment is not feasible for all types of policy authorities?" On the positive side, we can ask: "how would macroeconomic time series behave if there is a committed policy authority in place? how would the behavior be different if policy was made by an authority that could not commit?"

To benchmark the answers to these questions, we use a setting in which stochastic regime

[1] By describing the committed type as trustworthy, we mean that its announcements will correspond to its actions. In the current analysis, we do not focus on aspects of the announcement game, presuming that the conclusions reached in King et al. (2008) and Lu (2013) would carry over to the current setting.

transitions occur when type is known. With a committed type known to be in place, there is an interval of high but declining inflation, which gives way to zero average inflation for the duration of the regime. When it is known that the policy authority cannot commit, then there is high average inflation. Thus, the benchmark model can capture "great moderation" phenomena (1) with commitment and "inflation bias" phenomena (6) without commitment, as well as a specific cross-regime reason for time-varying response of inflation to energy price shocks frequently termed as "stabilization bias" from lack of commitment. Further, inflation expectations are low and stable under commitment but high and volatile without it, while only minimally influenced by prospective regime change in the calibrated economy that we study.[2]

Within a regime of unknown type, there is a richer interaction between policy actions, inflation history, expectations and real outcomes. To begin, a policy authority that cannot commit – which we sometimes call the alternative type – chooses its action understanding the consequences of observed inflation history and current inflation shocks for inflation expectations, although it does not believe it can manage expectations. In a setting where the private sector attaches a high likelihood to there being a commitment type in place, inflation expectations will lead the alternative type to be less aggressive (less biased) in its average choice and its response to observed inflation shocks. Accordingly, its optimizing behavior can lead to phenomenon (4), i.e., lengthy intervals of high real activity with gradually rising actual and expected inflation. As a history of positive inflation surprises leads to a decline in the likelihood that a committed type is in place, the alternative type will optimally choose to experience rising inflation accompanied by declining real activity (stagflation), i.e., phenomenon (5) will rise. Ultimately, its type will be revealed and it will simply experience high and volatile inflation with little effect on real economic activity.

[2]The stochastic regime transition with the policy authorities known to be committed is the same as the models of "loose committment" by Schaumburg and Tambalotti (2007) and Debortoli and Nunes (2014). The effect of prospective regime change is small in our model because we calibrate the model to have the average duration of a central banker's term to be 8 years.

In stressing that the optimal behavior of the alternative non-committed policy authority depends on the evolution of private sector inflation expectations, we have glossed over the fact that these are based in part on the optimal actions that a committed type would take given the same macroeconomic history. The optimal actions for the committed type do feature management of expectations, taking into account their effect on its own objective and on the behavior of an alternative non-committed type. The committed type understands that a lower private sector likelihood of its type being present – a lower reputation – worsens the leverage that its prospective policy actions have on current expected inflation. It also understands that the expected inflation affects the private sector's perceived optimal behavior of the alternative type, and in turn their inference about the policy authority's type from the random inflation outcome, i.e., the evolution of reputation.

The resulting optimal policy for a committed authority depends nonlinearly on its reputation. When the committed type has good reputation, its leverage on inflation expectations through the likelihood of its type being present extends to the perceived actions of the alternative type, as the optimal behavior of the non-committed type with good reputation is close to that of the committed type. The enhanced leverage on inflation expectations helps the committed type to stimulate output with an interval of high but declining inflation, similar to when the committed type is known to be in place. The private sector's learning is slow in this case. When the committed type has poor reputation, the private sector's inflation expectations are higher and more volatile not only because the likelihood of the committed type being present is low but also because the perceived optimal policies of the alternative non-committed authority are higher and more volatile. Such inflation expectations have an adverse effect on the committed type's output-inflation trade-off, forcing it to build up reputation rapidly through anti-inflation policies. Poorer reputation motivates more abrupt anti-inflation policies associated with larger output costs, which leads to phenomenon (2), i.e., anti-inflation policies, both gradual and abrupt, involving real costs.

Finally, the nonlinear effect of reputation on optimal committed policy also leads to time-varying responses of inflation to the relative price of energy, i.e., phenomenon (3), within a policy regime. In particular, the accommodation of inflation to an energy price shock is large when the committed type has poor reputation and is small when its reputation is good. As the private sector gradually learns that it is the committed type in place, the response of inflation to the energy price shock becomes smaller over the duration of the regime.

1.3. Related literature

This paper is closely related to the reputation literature on monetary policy. The early representative examples of this literature are Barro (1986), Backus and Driffill (1985*a*) and Backus and Driffill (1985*b*). They show that reputational force can discipline a discretionary policymaker to behave like a committed type, who mechanically adopts an exogenously given policy rule. Other papers in this literature put more emphasis on the optimal behavior of the committed policymaker. Cukierman and Liviatan (1991) and King et al. (2008) study how reputation concerns change the optimal committed monetary policy in a setting with the Lucas-Barro-Gordon Phillips curve. Hansen and McMahon (2016) and Xandri (2017) show the importance of signaling in monetary policy decisions, in which the "good" type of policymaker has a binary choice of signals. The most related paper is Lu et al. (2016), which studies optimal reputation building by the committed poilcymaker in a NK model with the non-committed policymaker mechanically following a given policy rule. Lu et al. (2016) finds that reputation building incentives make the committed policymaker less prone to simulate output with initially high but declining inflation and its inflation responses less accommodating to energy-price shocks. In the current paper, we treat both types of policymakers (committed or non-committed) as strategic players and solve for their optimal inflation policies in a single recursive optimization problem. We find that taking into account the optimal responses of a non-committed policymaker significantly changes the reputation building incentives of the committed policymaker, making the effect of reputation on optimal committed policy highly nonlinear.

In this paper, the rich dynamics of reputation is governed by the private sector's learning about the type of the policy authority in place. Our model thus belongs to the vast learning literature on monetary policy,[3] which is well known to be more consistent with basic facts about measured expectations and forecasting errors than the standard rational expectation (RE) approach. Most papers in the literature assume that the agents learn using a misspecified model of the economy. Our approach differs in that our private sector possesses prefect knowledge about the economic model (i.e., the mapping from the states to the policies), yet incomplete information of the type of policy authority. These two approaches are complementary in capturing different uncertainty regarding the fundamentals of the economy. Our approach is more appropriate for studying an economy in which the optimal monetary policy is well understood given the central banker's preference and commitment technology. Adopting a similar approach, Matthes (2015) estimates a model in which the private sector uses Bayes' law on a rolling data sample to discriminate between two models of monetary policymaking. He finds that the private sector increasingly believed that the monetary policy was set with commitment during the Volcker disinflation. Our model differs from his in that the optimal monetary policy – with or without commitment – takes into account the private sector's learning, and the private sector understands such sophistication of the optimal monetary policy.

Finally, this paper is related to the recursive contracts literature since the equilibrium of the model is the solution to a recursive optimal policy problem for the committed type. It is now standard to analyze optimal policy under commitment using methods from the recursive contracts literature, with the policy authority (the principal) selecting macroeconomic outcomes that can be implemented in a particular market structure by the decentralized

[3] See Evans and Honkapohja (2003), Evans and Honkapohja (2008), Woodford (2013) and Eusepi and Preston (2018) for the surveys of this literature.

choices of households and firms.[4] Relative to a standard dynamic principal agent problem, our model possesses several new elements worth stressing. First, the principal (the committed type) and the agent (the non-committed type) are never simultaneously present. The type of monetary authority that is present is governed by an exogenous probability structure, in which publicly observed replacement event yields a randomly drawn type of new policy authority. Second, the principal cares only about the value of the objective if he is present, but that of the agent. Third, due to the fact that nature determines whether the principal or agent is present, contract design involves only an incentive compatibility constraint for the agent and not a participation constraint. Fourth, the rewards to either the principal or the agent depend on the beliefs of a third party (the private sector), specifically the likelihood that each type of authority will be present next period and the actions that each would take given the future state of the economy.

2. The Economy

Consider an economy in which a policymaker designs and implements policies that are payoff-relevant to a private sector composed by forward-looking agents. The private agents are unsure whether the poilcymaker can commit or not. In order words, when the private agents make forward-looking decisions, they take into account the possibility that an announced policy plan may not be executed.

2.1. Policymakers: types and actions

The policymaker in the economy is a central banker who is responsible for the inflation rates. A central banker can serve for multiple periods, with the term length determined randomly by a constant replacement probability q each period. A replacement is not forecastable but is a publicly observed event if it occurs.

[4] Examples can be found but are not limited in: Khan et al. (2003), Golosov et al. (2016) and Marcet and Marimon (2017).

We study the optimal inflation policy of a central banker who maximizes the payoffs within his own term. The central banker can be one of the two types: the *committed* type or the *alternative* type.

The committed type of central banker is capable of commitment. He chooses the optimal inflation plan when he first takes office and commits to the plan for all subsequent periods, conditional on his holding office. The predetermined optimal plan specifies the committed type's inflation policy in each period t, denoted by a_t, contingent on the realization of shocks. The committed type has a time discount factor β_1 and a momentary payoff:

$$u_t = -\frac{1}{2}[\pi_t^2 + h_1(x_t - x^*)^2], \tag{1}$$

where π_t is the inflation rate, x_t is the output gap, and $x^* > 0$ is the output gap target.

The alternative type of central banker cannot commit so that he chooses the optimal inflation policy, denoted by α_t, on a period-by-period basis. We assume that the alternative type is myopic and his momentary payoff is

$$\upsilon_t = -\frac{1}{2}[\pi_t^2 + h_2(x_t - x^*)^2]. \tag{2}$$

To help exposition later, we assume that the central banker announces his inflation plan in advance. The committed type will announce a_t and honor the announcement by his implemented policy. The alternative type will also announce a_t but will implement α_t [5]

2.2. Private sector: payoff and information

A private sector is composed by all agents in the economy but the policymaker. There are two key characteristics of the private agents: 1) none of them has strategic power (they are

[5] In our prior work, King et al. (2008) and Lu (2013), we studied the signaling equilibrium of a model in which the alternative type chooses optimal announcement strategy. We found that both types of policy authorities would make the same announcement and that the announcement would be the optimal policy for the committed type. Our present assumption is in line with these findings.

atomistic); 2) they are forward-looking. A reduced-form aggregate outcome of the interaction among the private agents is represented by a standard New Keynesian (NK) Phillips curve

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + \varsigma_t, \tag{3}$$

where β is the private sector's time discount factor, $E_t \pi_{t+1}$ is the expected future inflation, and the cost-push shock ς_t follows an exogenous Markov chain process:

$$\Pr\left(\varsigma_{t+1} = s | \varsigma_t = \sigma\right) = \delta\left(s, \sigma\right). \tag{4}$$

The private sector does not observe the central banker's type or his inflation policy, a_t or α_t. Yet, it observes an inflation rate π_t that is a random outcome of the inflation policy. In particular:

$$\pi_t = \begin{cases} a_t + \varepsilon_t & \text{under the committed type} \\ \alpha_t + \varepsilon_t & \text{under the alternative type} \end{cases}, \tag{5}$$

where ε_t is an i.i.d. implementation error with mean zero, variance σ_ε^2, and a bell-shaped distribution that peaks at zero.[6] We interpret this random inflation error as a reduced-form representation for all unforeseeable factors that affect the inflation rate beyond the monetary policy. There is ample evidence that realized inflation rates often miss the inflation targets, with examples including Roger and Stone (2005) and Mishkin and Schmidt-Hebbel (2007).

2.3. Reputation, credibility, and expected inflation

The realized inflation rate is thus a noisy signal of the unobserved policy, or equivalently, the type of the incumbent central banker. Denote by ρ_t the private sector's assessment (as of the start of period t) of the probability that the incumbent central banker is the committed

[6] A similar structure with implementation error can be found in Cukierman and Meltzer (1986), Faust and Svensson (2001), Atkeson and Kehoe (2006), etc.

type. After observing π_t, the private sector updates ρ_t according to Bayes' rule:

$$\rho_{t+1} = \frac{\rho_t f(\pi_t|a_t)}{\rho_t f(\pi_t|a_t) + (1-\rho_t) f(\pi_t|\alpha_t)}, \tag{6}$$

where $f(\pi|a)$ denotes the probability of observing π, conditional on the underlying policy's being a. We refer to ρ as the *reputation* of the central banker.

The central banker's reputation ρ determines the extent to which the announced policy plans can affect the private sector's expected inflation:

$$e_t \equiv \beta E_t \pi_{t+1} = \beta \left\{ \begin{array}{c} (1-q)\left[\rho_{t+1} E_t(a_{t+1}) + (1-\rho_{t+1}) E_t(\alpha_{t+1})\right] \\ +q\left[\rho_0 E_t(a_0) + (1-\rho_0) E_t(\alpha_0)\right] \end{array} \right\}. \tag{7}$$

In this expression, the right hand side is a probability-weighted average of the expected inflation that will take place if the incumbent central banker continues his term and that if the incumbent central banker is replaced. If the current central banker continues his term, the expected inflation will be a_{t+1} conditional on the incumbent central banker being the committed type, and α_{t+1} conditional on him being the alternative type, so that $E_t(a_{t+1})$ is weighted by ρ_{t+1} and $E_t(\alpha_{t+1})$ is weighted by $1-\rho_{t+1}$. Turning to the case of a replacement, the newly appointed central banker will reoptimize and set a_0 if he is the committed type and α_0 if he is the alternative type. We specify the initial reputation of the newly appointed central banker as follows:

$$\rho_0 = \rho(\tau, \omega, \rho_{t+1}) = \tau + \omega(\rho_{t+1} - \tau).$$

One can interpret ω as the probability of the new central banker being of the same type as his predecessor. With the complementary probability, $1-\omega$, the type of the new central banker will be an unconditional draw that has τ probability to be the committed type. Note that the appropriate reputation measure ρ or $\tau + \omega(\rho - \tau)$ captures the *credibility* of the announced policy plan a_{t+1} or $a_{0.}$

2.4. Timing of events

Within a period t, events take place in the following order. A public replacement of the incumbent central banker either occurs or not. The cost-push shock ς_t hits. An inflation plan will be announced if the central banker is a new comer. Otherwise, there will be no announcement. Then the central banker implements a hidden inflation policy, a_t or α_t, depending on his type. This hidden policy results in an observable inflation outcome π_t. The private sector forms expectations about inflation in the next period, e_t. Finally, the output gap x_t is determined by the Phillips curve.

3. Macro Equilibrium

The economy just described consists of a private sector and a central banker who can be one of the two types, but whose actions do not directly reveal his type: we thus are essentially considering a dynamic game with imperfect information. This section defines equilibrium in this dynamic game and analyzes the strategic interaction between the two types of central bankers, in light of the learning behavior of private agents. In closing the section, we describe a principal-agent interpretation of the dynamic game.

3.1. Public Equilibria

Define the publicly observable history $h_t = \{h_{t-1}, \pi_{t-1}, \varsigma_t\}$ as the collection of all past realizations of inflation rates and cost-push shocks. We restrict our attention to equilibria in which all strategies are conditional only on the public history, i.e., "public strategies."

Such a restriction is innocuous in our equilibrium analysis because: 1) the private sector's strategy has to be public since h_t is the private sector's information set; 2) the committed type's policy has to be public since it follows the announced policy plan, which needs to be verifiable by the private sector;[7] 3) given all the other player's strategies are public, it is also optimal for the alternative type to choose public strategies (Mailath and Samuelson (2006)).

3.2. *Perfect Bayesian*

We further require the equilibrium in this imperfect information game to be perfect Bayesian. That is, the beliefs of the private sector are consistent and the strategies of the two types of central bankers satisfy sequential rationality.

3.2.1. *Consistent beliefs*

The expected inflation in (7) summarizes all the relevant beliefs of the private sector: how likely is the incumbent central banker the committed type and what are the expected inflation policies under the committed and the alternative types of the incumbent and the new central bankers. The posterior belief of the incumbent central banker's type follows the *Bayes' rule* as in (6). The inflation policy of the incumbent committed type follows the announced policy plan is thus known to the private sector. The other inflation policies, however, are not announced so that the private sector needs to make conjectures on what they will be. In equilibrium, we impose *rational expectations* in the sense that the conjectured policies coincide with the equilibrium policies.

The consistency restriction on the beliefs of private sector therefore implies that the posterior belief ρ_{t+1} in (6) is a function of (h_t, π_t) and the expected inflation e_t in (7) can be written as

$$e(h_t, \pi_t) = \beta \left\{ \begin{array}{c} (1-q)\left[\rho(h_t, \pi_t) E_t\left[a(h_{t+1})\right] + (1-\rho(h_t, \pi_t)) E_t\left[\alpha(h_{t+1})\right]\right] \\ +q\left[\rho_0(h_t, \pi_t) E_t\left(a_0(h_{t+1})\right) + (1-\rho_0(h_t, \pi_t)) E_t\left(\alpha_0(h_{t+1})\right)\right] \end{array} \right\}, \quad (8)$$

where $E_t(\cdot) = \sum_{\varsigma_{t+1}} \delta(\varsigma_{t+1}, \varsigma_t)(\cdot)$.

[7] We do not allow mixed strategy in announcement, which can be an interesting extension of the current model.

3.2.2. Sequential rationality of the alternative type

Although the central banker of the alternative type implements the inflation policy before the private sector's actions, his inflation policy has no strategic power on the private sector's expectation because neither does his policy follow the announced one nor is it observable.[8] The alternative type of the central banker therefore maximizes his momentary payoff (2), taking as given the private sector's expected inflation $e(h_t, \pi_t)$, the effect of his policy α on the distribution of the inflation rate π, and the Phillips curve (3).

With moderate conditions on $e(h_t, \pi_t)$ to ensure the interior solution to the alternative type's optimization, a sequential rational inflation policy α_t needs to satisfy the first order condition:

$$\int \upsilon\left(\pi_t, e\left(h_t, \pi_t\right), \varsigma_t\right) f_\alpha\left(\pi_t | \alpha_t\right) d\pi_t = 0, \tag{9}$$

where $\upsilon\left(\pi_t, e_t, \varsigma_t\right)$ is the momentary objective with x_t replaced by $\left(\pi_t - e_t - \varsigma_t\right)/\kappa$, and $f_\alpha\left(\pi_t | \alpha_t\right)$ is the first-order derivative of the density function $f\left(\pi | \alpha\right)$.

A special case with normal density: If we assume normal density function for $f(\pi|\alpha)$, we can further solve for the solution to the first order condition (9):

$$\alpha_t = \frac{h_2}{\kappa^2 + h_2}[E_t^- e_t + \varsigma_t + \kappa x^* + HOT]. \tag{10}$$

where $E_t^- e_t = \int e\left(h_t, \pi_t\right) f\left(\pi_t | \alpha_t\right) d\pi_t$ and HOT is the higher-order term that collects second or higher moments:

$$HOT = -\frac{cov(e_t, \pi_t - \alpha_t)}{var(\pi_t)}\left(\varsigma_t + \kappa x_t^* - \alpha_t\right) - \frac{1}{2}\frac{cov(e_t^2, \pi_t - \alpha_t)}{var(\pi_t)} + \frac{cov(e_t, (\pi_t - \alpha_t)^2)}{var(\pi_t)}.$$

[8] Because it is impossible for private agents to adjust their belief based on some unobservable change of alternative type's inflation policy. The situation is analogous to Osborne and Rubinstein (1994)'s discussion of the Stackelberg game.

The solution in (10) highlights how the optimal inflation policy of the alternative type α_t depends on the private sector's expectation function $e(h_t, \pi_t)$.

If e_t is independent of π_t, the solution in (10) is the one with both inflation bias and stabilization bias derived in a pure discretion model. If e_t is responsive to π_t, there are two additional channels that affect α_t.

First, when e_t works to smooth the effect of inflation shock on the output, the optimal alternative policy α_t is lower than the pure discretionary rule. In an extreme case where the smoothing effect is perfect and the output is always at its targeted level x^*, the optimal alternative policy is zero. This effect is captured by the terms $E_t^- e_t + \varsigma_t + \kappa x^*$ in (10).

Second, e_t can have different sensitivity to π_t conditional on the realization of π_t. The optimal alternative policy determines the mean of the realized inflations and it is chosen to balance the inflation risk and the implied risk on output. Therefore, if e_t is better at smoothing the inflation shock at higher values of π_t, the optimal alternative policy will be higher so as to put more probability weight on these low-output-risk inflation rates. Conversely, the optimal alternative policy tends to be lower if e_t is more sensitive to π_t at lower values of π_t. This effect is captured by the HOT term in (10).

3.2.3. Sequential rationality of the committed type

The central banker of the committed type, in contrast to the alternative type, decides all his future state-contingent policies at period 0 (the beginning of his term) and commits to them. Through announcement, his committed state-contingent policy plan is public information and therefore has strategic power on both the private sector's expectation and the alternative type's inflation policies. In particular, the strategy of the committed type is *sequentially rational* if it maximizes his expected payoff at period 0,

$$E_0\{\sum_{t=0}^{\infty} \beta_1^t (1-q)^t \int u(\pi_t, e(h_t, \pi_t), \varsigma_t) f(\pi_t|a_t) d\pi_t\}, \tag{11}$$

where $u(\pi_t, e_t, \varsigma_t)$ is the momentary objective of the committed type with x_t replaced by $(\pi_t - e_t - \varsigma_t)/\kappa$. The committed type takes into account that the private sector's expectation $e(h_t, \pi_t)$ is formed based on a *consistent* belief system (8), through which both the committed and the alternative policies affect how e_t depends on π_t, and that the alternative type's period-t inflation policy is sequentially rational, i.e., it satisfies (9), and hence is an optimal response to the period-t private sector's expected inflation $e(h_t, \pi_t)$.

3.3. A principal-agent interpretation

Constructing the Public Perfect Bayesian equilibrium in this game can be usefully viewed as solving a principal-agent problem. The central banker of the committed type is the principal. He chooses the future actions of his own and the future actions for the two agents – the private sector and the central banker of the alternative type. In choosing the future actions of the two agents, the principal has to respect two incentive compatibility constraints, one is the rational expectation constraint for the private sector (8) and the other is the sequential rationality constraint for the alternative type (9). Our further analysis will be based on this principal-agent interpretation.

4. Recursive Framework

The optimal policy problem for the central banker of the committed type (the principal as in the principal-agent interpretation) is to maximize the objective (11) subject to rational private sector's expectation (8) and the sequential rationality of the alternative type (9). Under the first order approach, we can form a Lagrangian by attaching multipliers γ_t to the expectation constraint and ϕ_t to the sequential rationality constraint. Then, as shown in the appendix, we can rewrite the infinite-horizon maximization problem in a recursive form with state variables ρ_t and ς_t and a pseudo state variable η_t,[9] following the research path

[9] The pseudo state variable is proportional to the lagged multiplier γ_{t-1}, which is another manner in which the history dependence is frequently incorportated.

laid out by Kydland and Prescott (1980), Marcet and Marimon (2017) and others.[10]

$$W\left(\rho_t, \eta_t, \varsigma_t\right) = \min_{\phi_t, \gamma_t(\pi_t)} \max_{a_t, \alpha_t, e_t(\pi_t)} E_t^- \left\{ w_t + \beta_1 (1-q) E_t W\left(\rho_{t+1}, \eta_{t+1}, \varsigma_{t+1}\right)\right\}, \tag{12}$$

where $E_t^-(\cdot) = \int (\cdot) f(\pi_t | a_t)\, d\pi_t,\ E_t(\cdot) = \sum_{\varsigma_{t+1}} \delta(\varsigma_{t+1}, \varsigma_t)(\cdot)$,

$$w_t = u\left(\pi_t, e\left(\pi_t\right), \varsigma_t\right) + \gamma_t\left(\pi_t\right) e_t(\pi_t) \tag{13}$$

$$- (1-q)\, \eta_t \left[\rho_t a_t + (1-\rho_t)\alpha_t\right] \tag{14}$$

$$-q\eta_t \left[\left(\left(1-\omega\right)\tau + \omega\rho_t\right) a_0 + \left(1 - \left(1-\omega\right)\tau - \omega\rho_t\right)\alpha_0\right] \tag{15}$$

$$+\phi_t \frac{f_\alpha\left(\pi_t|\alpha_t\right)}{f\left(\pi_t|a_t\right)} \upsilon\left(\pi_t, e\left(\pi_t\right), \varsigma_t\right), \tag{16}$$

subject to state evolution equations for the cost-push shock (4) and the reputation state (6), and

$$\eta_{t+1} = \frac{\beta}{\beta_1 (1-q)} \gamma_t, \text{ with } \eta_0 = 0. \tag{17}$$

Notice that ϕ_t, a_t and α_t are chosen before the realization of π_t, whereas γ_t and e_t are functions of π_t because the private sector forms expectation after observing π_t. Within this approach, our task is thus to determine how equilibrium outcomes such as a depend on the state vectors $s_t = (\rho_t, \eta_t, \varsigma_t)$, how equilibrium outcomes such as e depend on (π_t, s_t) and how the pseudo state η evolves with (π_t, s_t).

4.1. Familiar elements of policy design

It is useful to consider the elements of policy design in (12) that are familiar from textbook analyses such as those of Woodford (2003), Walsh (2010), and Gali (2015), and from earlier work in the literature such as Schaumburg and Tambalotti (2007) and Lu et al. (2016).

[10]See also Chang (1998) and Phelan and Stacchetti (2001).

Full commitment benchmark Optimal monetary policy with non-stochastic inflation ($\sigma_\varepsilon = 0$), no regime changes ($q = 0$) and perfect credibility ($\rho = 1$) is well understood to be history dependent, i.e., there is an interval of startup inflation. This feature is captured in the recursive framework by the dependence of the policy a on the pseudo state variable η_t with startup inflation being transitional dynamics from the initial condition $\eta_0 = 0$. The term $-\eta_t a_t$ in (14) captures the effect of past commitments: a greater value of η makes it less desirable to raise current inflation from 0. Along the transitional path, as the pseudo state variable η rises over time towards its positive steady state value, the policy action a declines and, ultimately, converges to zero long-run inflation.

Such a path of initially high but declining inflation highlights that the optimal policy involves expectations management. In each period of the transition, the expected inflation is engineered to be lower than the current inflation $e_t < \pi_t$, so that there is a stimulation of real activity brought by the optimal policy. Taking first order condition of e_t of the recursive program (12),

$$\gamma_t = -u_{e,t} = -\frac{h_1}{\kappa}\left(x_t - x^*\right), \tag{18}$$

reveals that the choice of e_t balances the beneficial effect of lower expected inflation ($-u_{e,t}$), which works to stimulate real activity at a given inflation rate as long as the output gap $x_t = \frac{1}{\kappa}\left(\pi_t - e_t - \varsigma_t\right) < x^*$, with the fact that future inflation must be lower given rational expectations (a higher γ_t increases η_{t+1} and in turn lower a_{t+1}).

Stochastic inflation With stochastic inflation ($\sigma_\varepsilon > 0$), expectation management takes on a heightened role. For example, after a positive inflation shock ($\pi - a = \varepsilon > 0$), output will be unexpectedly high and it is optimal to smooth out such a beneficial surprise over time by raising expected inflation. Technically, a positive inflation shock in period t produces a contemporaneous decline in the multiplier γ_t, which translates into a decline in the pseudo state variable η_{t+1} in the following period and an increase in the optimal policy a_{t+1}, thus ratifying the efficient movement in expected inflation $\Delta e_t > 0$. Just as in the transitional dynamics, the efficient response to the unexpected positive inflation shock is to produce a temporary interval of high but declining inflation.

Prospective regime change With prospective regime change ($q > 0$), there are elements of policy design stressed by Schaumburg and Tambalotti (2007): the future is discounted more heavily in the objective (12) and the optimal policies of a successor appear in forward-looking constraint, i.e., in (15) within the recursive program. Since a new regime starts without any prior commitment ($\eta_0 = 0$), an equilibrium requires that optimal policies $a(\rho, 0, \varsigma; a_0, \alpha_0) = a_0$ and $\alpha(\rho, 0, \varsigma; a_{0,}\alpha_0) = \alpha_0$. Our computations employ a fixed point recipe: we conjecture $a_0(\rho, \varsigma)$ and $\alpha(\rho, \varsigma)$, solve for optimal a and α, the iterate until converge.

Evolving reputation With endogenous reputation, there is a state variable ρ not present in the standard NK framework, whose evolution is influenced by the committed type's action a_t, but also by inflation outcome π_t and the alternative type's action α_t. That is, Bayesian learning (6) specifies that $\rho_{t+1} = b(\pi_t, a_t, \alpha_t, \rho_t)$. Accordingly, an additional element of policy choice is the implications for reputation evolution, as explored in Lu et al. (2016) in a setting where there was a mechanical alternative type of central banker that adopted a rule $\alpha_t = \alpha(\varsigma_t)$. While it is possible to calculate exact linear solutions for models with stochastic inflation and prospective regime change, the nonlinearity of Bayes' rule means that nonlinear numerical methods must be used in settings with reputation evolution.

4.2. Policy design with incentive compatibility

The novel element of this study is that α_t is chosen optimally by the alternative type of central banker. This feature is captured in the recursive framework by including (16) as the incentive compatibility constraint for the alternative type's policy. Notice that this incentive compatibility constraint differs from the original sequential rationality constraint (9) by a probability transformation $\frac{f_\alpha(\pi_t|\alpha_t)}{f(\pi_t|a_t)}$. It is because the alternative type's policy is chosen prior to the realization of inflation so the sequential rationality constraint takes into account that the random inflation outcome has mean equal to α_t, whereas the recursive program is conditional on the central banker being the committed type so that the random inflation outcome has mean equal to a_t.

4.2.1. Conditionally optimal expectation function

Including the incentive compatibility constraint (16) modifies the first order condition of e_t to

$$\gamma_t = -u_{e,t} - \upsilon_{e,t}\phi_t \frac{f_\alpha\left(\pi_t|\alpha_t\right)}{f\left(\pi_t|a_t\right)} \tag{19}$$

where $u_{e,t}$ is the same as in (18) and $\upsilon_{e,t} = \left(h_2/\kappa\right)\left(x_t - x^*\right)$ is the derivative of the alternative type's momentary objective υ_t with respect to e_t. In other words, expectation management needs to take into account the effect of e_t on the optimal policy of the alternative type.

To see it more explicitly, we combine the first order condition with the requirement that the beliefs of private sector are consistent (8) to determine a conditionally optimal expectation function. For the sake of easier exposition, let us define

$$\begin{aligned} &\widehat{M}\left(\rho_{t+1}, \eta_{t+1}, \varsigma_t\right) \\ \equiv\ &\sum_{\varsigma_{t+1}} \delta\left(\varsigma_{t+1}, \varsigma_t\right) \left[\begin{array}{c} (1-q)\left[\rho_{t+1}a_{t+1} + (1-\rho_{t+1})\alpha_{t+1}\right] \\ +q\left[\left((1-\omega)\tau + \omega\rho_{t+1}\right)a_0 + \left(1-(1-\omega)\tau - \omega\rho_{t+1}\right)\alpha_0\right] \end{array} \right]. \end{aligned}$$

so that the consistency of beliefs (8) may be written as:

$$e_t = \beta\widehat{M}\left(\rho_{t+1}, \eta_{t+1}, \varsigma_t\right). \tag{20}$$

This representation (20) reveals that the private sector's expectations rationally anticipate the next-period equilibrium policies as functions of the future state variables $\left(\rho_{t+1}, \eta_{t+1}, \varsigma_{t+1}\right)$. ρ_{t+1} is the private sector's updated beliefs that the incumbent central banker is a committed type, which is determined by $(\pi_t, a_t, \alpha_t, \rho_t)$ according to the Bayes' rule. The pseudo-state variable η_{t+1}, according to (17), is proportional to γ_t, which is in turn determined in (19) by $\pi_t, e_t, \varsigma_t, \phi_t, a_t$ and α_t. The next-period cost-push shock ς_{t+1} evolves exogenously with transitional probability depending solely on ς_t.

Hence, there is a conditionally optimal expectation function that is implicitly determined by equations (19), (17), and (20). Our analytical and computational strategy is to find the triplet $\left(e_t, \gamma_t, \eta_{t+1}\right)$ that satisfies these three equations for each π_t given $(\phi_t, a_t, \alpha_t, \rho_t, \varsigma_t)$. From the perspective of managing expectation, given the state variables (ρ_t, ς_t), the optimal expectation function $e\left(\pi_t\right)$ conditional on the current policies (a_t, α_t) varies with the choice of ϕ_t. To emphasize that dependence, we write the conditionally optimal expectation function as

$$e\left(\pi_t | \phi_t; a_t, \alpha_t, \rho_t, \varsigma_t\right).$$

It turns out that the choice of ϕ_t is restricted by the second order condition of e_t in the recursive problem (12):

$$-\frac{h_1}{\kappa^2} - \frac{h_2}{\kappa^2}\phi_t \frac{f_\alpha\left(\pi_t|\alpha_t\right)}{f\left(\pi_t|a_t\right)} < 0, \text{ for } \forall \pi_t$$

This condition translates to an admissible set for ϕ_t as follows:

$$\phi_t \in \Phi(a_t, \alpha_t) \text{ where } \Phi(a_t, \alpha_t) \equiv \begin{cases} \left[\underline{\phi}_t, 0\right] & \text{if } \alpha_t < a_t; \\ \left[0, \overline{\phi}_t\right] & \text{if } \alpha_t > a_t; \end{cases} \tag{21}$$

where $\underline{\phi}_t$ and $\overline{\phi}_t$ are functions of (a_t, α_t).[11] Therefore, there is a bounded set of expectation functions $e\left(\pi_t|\phi_t; a_t, \alpha_t, \rho_t, \varsigma_t\right)$ from which the committed type may choose.

[11] $\underline{\phi}_t = \overline{\phi}_t \equiv -\frac{h_1}{h_2} \frac{a_t - \alpha_t}{\exp\left(\frac{1}{2\sigma^2}\left[(\widehat{\pi}_t - a_t)^2 - (\widehat{\pi}_t - \alpha_t)^2\right]\right)}$ with $\widehat{\pi}_t = \alpha_t + \frac{\sigma^2}{a_t - \alpha_t}$.

4.2.2. Incentive compatible set of α_t

The section above shows how the conditionally optimal expectation function depends on the current policies (a_t, α_t). Yet the sequential rationality constraint of the alternative type (9) imposes the dependence of the optimal α_t on the expectation function. Technically, the multiplier ϕ_t on the sequential rationality constraint has to be adjusted until the constraint is satisfied, if a particular a_t, α_t pair is to be incentive compatible. Define

$$I\left(\phi_t | a_t, \alpha_t\right) \equiv \int \upsilon\left(\pi_t, e\left(\pi_t | \phi_t; a_t, \alpha_t\right), \varsigma_t\right) f_\alpha\left(\pi_t | \alpha_t\right) d\pi_t \tag{22}$$

as the left hand side of the constraint (9). We have the following result:

Lemma 1 *$I\left(\phi_t | a_t, \alpha_t\right)$ is strictly increasing in ϕ_t.*

The lemma implies that there is a unique ϕ_t to make $I\left(\phi_t | a_t, \alpha_t\right) = 0$, or equivalently, there is a unique expectation function that the committed type may use to make a particular a_t, α_t pair incentive compatible. Furthermore, as shown in the earlier section, the conditional optimality of e_t constrains the choice of ϕ_t to an admissible set $\Phi(a_t, \alpha_t)$ as in (21). It is possible that the unique ϕ_t that makes $I\left(\phi_t | a_t, \alpha_t\right) = 0$ is outside the admissible set, in which case the particular a_t, α_t is not incentive compatible. We use such a criterion to identify the set of incentive compatible α_t given a particular committed policy a_t.

Definition 2 *Given (ρ_t, ς_t) and a_t, the incentive compatible set of α_t is defined as:*

$$IC\left(a_t; \rho_t, \varsigma_t\right) \equiv \left\{\alpha_t : \exists \hat{\phi} \in \Phi(a_t, \alpha_t) \text{ s.t. } I\left(\hat{\phi} | a_t, \alpha_t\right) = 0\right\}.$$

4.2.3. Summary

To summarize, the first and second order conditions of e_t, the rational expectation constraint, and the sequential rationality constraint of the alternative type (the latter two are the first order conditions of γ_t and ϕ_t, respectively) identify a set of α_t that can be incentive compatible given a particular a_t. Each incentive compatible α_t with a given a_t is associated with one conditionally optimal expectation function $e\left(\pi_t|\hat{\phi}(a_t,\alpha_t);a_t,\alpha_t,\rho_t,\varsigma_t\right)$ where $\hat{\phi}(a_t,\alpha_t)$ belongs to the set $\Phi(a_t,\alpha_t)$ and makes $I\left(\hat{\phi}(a_t,\alpha_t)|a_t,\alpha_t\right)=0$.

Intuitively, the value of ϕ_t summarizes how the incentive compatibility of α_t affects expectation management. On one hand, having α_t respond to the choice of expectation function helps the committed type to better smooth shocks and manage reputation. Technically, α_t enters the term $-\eta_t(1-\rho_t)\alpha_t$ in (14) that captures the effect of past commitments, and the state evolution equation (6) for the next-period reputation ρ_{t+1}. On the other hand, the choice of a particular expectation function constrains future inflation policies given rational expectations. It is captured by the effect of ϕ_t on η_{t+1} through the first order condition of e_t (19) and the state evolution equation for η_{t+1} (17).

After identifying the incentive compatible (IC) set conditional on a given (ρ_t,ς_t), we then use the $e\left(\pi_t|\hat{\phi}(a_t,\alpha_t)\right)$ and $\gamma\left(\pi_t|\hat{\phi}(a_t,\alpha_t)\right)$ associated with each (a_t,α_t) pair to evaluate the committed type's payoff as in the RHS of recursive program (12). Notice that evaluating the committed type's payoff is the only place where the state η_t enters the program. Define the state of the program by $s_t\equiv(\rho_t,\eta_t,\varsigma_t)$. The (a_t,α_t) pair that yields the highest payoff for the state s_t determines the equilibrium inflation policies $a(s_t)$, $\alpha(s_t)$ as well as the expectation function $e(\pi_t,s_t)$ and the multiplier function $\gamma(\pi_t,s_t)$.

Note that in the above analysis and in the computational approach, we have avoided utilizing the first order conditions of a_t and α_t. We did so because, ex ante, we knew little about the nature of the incentive compatible set over which the committed type can optimize. The literature on contract design has taught us that the incentive compatible set can be complicated, even sometimes invalidating the first order condition approach. We thus adopted a strategy of (i) optimizing over (a,α) pairs in the IC set; and (ii) performing various checks on our derivation of the IC set.

5. Equilibrium Policies

In this section, we use a calibration model approach to study how the equilibrium dynamics evolves along the transitional path following a replacement of the central banker, and how the equilibrium dynamics respond to the cost-push shock and the implementation error (interpreted as the unexpected demand shock). The variables in interest in the equilibrium dynamics are the inflation policy, the output, the reputation and the expected inflation. The equilibrium dynamics is conditional on the underlying type of central banker that is in place and his initial reputation. We are going to underline the importance of the incentive compatibility of the alternative inflation policy in producing these equilibrium dynamics as it is at the center of this research.

Table 1 summarizes the parameter values in our benchmark calibration.

β, β_1	Discount factor	0.995
q	Replacement probability	0.03
h_1, h_2	Output weight	0.017
x^*	Output target	0.05
κ	PC output slope	0.17
σ_ε	Std of implementation error	0.5%
σ_ξ	Std of cost-push innovation	0.5%
δ	Persistence of cost-push shock	0.9
ρ_0	Initial reputation after replacement	$1\% + 0.5\rho_{-1}$

Table 1

Most parameters in this calibration share the common values with the baseline calibration in Lu et al. (2016) as it is a fairly standard one for NK models. The same time discount factor is assumed for both the committed central banker and the private sector. It implies a steady-state interest rate of about 2% annually. The replacement probability q is set to 0.03 so that the average tenure of a central banker in our infinite horizon model matches the 8 years finite horizon in the model of Lu et al. (2016). $\{h, x^*, \kappa, \beta\}$ imply an annual inflation bias equal to 2%, which is admittedly too small to match the inflation experience of the US and other countries. But since this paper focus on capturing some "stylized features" of the U.S. inflation experience rather than matching it quantitatively, this magnitude of inflation bias is sufficient and it facilitate exposition and comparison of the numerical results with the ones in Lu et al. (2016) where the alternative type of the central banker is assumed to follow a mechanical rule.

As in Lu et al. (2016), the parameters $\{h, x^*, \kappa\}$ are chosen to be consistent with the microfoundation of the central banker's objective as a second-order expansion of the representative consumer's utility (Gali (2015)) and the estimated Phillips curve using a marginal cost proxy (Gali and Gertler (1999)). The implied values of the structural parameters of the underlying economy from $\{h, x^*, \kappa, \beta\}$ are:[12] the elasticity of marginal cost with respect to the output ($A = 2$); the demand elasticity ($\epsilon = 10$); and the probability that a firm is able to reoptimize its nominal price each quarter ($1 - \theta = 0.25$). Finally, we assume the implementation error follows a normal distribution with a standard deviation σ_ε equivalent to 1% annually (as opposed to 2% annually in Lu et al. (2016)). The empirically relevant range of the implementation error's standard deviation is $[0.8\%, 2\%]$ as documented by Mishkin and Schmidt-Hebbel (2007) and Roger and Stone (2005). In this exercise, we choose σ_ε near the lower bound of its empirical range to maximize the committed central banker's incentive to build reputation. The cost-push shock is modeled as a Markov chain: $\varsigma_t = \varsigma_{t-1}$ with probability δ and $\varsigma_t = \xi_t$ with probability $1 - \delta$, where the innovation ξ_t is uniformly distributed over $[-\overline{\varsigma}, \overline{\varsigma}]$ with the standard deviation σ_ξ. The parameter values are the same

[12] The details of the mapping between $\{h, x^*, \kappa, \beta\}$ and $\{A, \epsilon, \theta\}$ can be found in the appendix B of Lu et al. (2016).

as in Lu et al. (2016). Finally, after replacement, the newly appointed central banker has a 50% chance to be the same type as his predecessor and a 50% chance to be a committed type with probability 2%. This initial condition is chosen to motivate investment in reputation.

5.1. Transitional dynamics

When an incumbent central banker is publicly replaced, the newcomer reoptimizes regardless of his type. This subsection presents transitional dynamics of the new regime without any preexisting commitments, i.e., $\eta_0 = 0$. We assume that the realized implementation errors and cost-push shocks are zero in all periods to focus on the transitional dynamics and leave the discussions of the responses to shocks to the next subsection. As a result, the realized inflation rates equal the equilibrium inflation policies implemented by the central banker in place. Since the private sector does not observe the realized shocks, they still need to learn about the central banker's type, but their belief of the central banker being the committed type will steadily grow if the central banker is the committed type and will steadily decline otherwise. The contrasting transitional paths of the private sector's belief, which is a key state variable, lead to different transitional dynamics of other equilibrium variables including the inflation policies, the output, and the expected inflation. Finally, since our model is nonlinear, the value of the initial reputation ρ_0 also matters for the transitional dynamics.

5.1.1. New regime with a committed central banker

Figure 1 plots the entire path of the optimal *committed* inflation policy (also the realized inflation rate) in the first 4 years after a committed central banker takes the office, together with the implied output and reputation. The transitional dynamics with 6 different levels of initial reputation are plotted: 1('*'), .75('Δ'), .5('$\Diamond$'), .2('∇'), .1('+'), 0('$\bigcirc$'). The inflation rate is plotted at its annualized level and the output is in percent deviation from its distorted steady state.

The cases with $\rho_0 = 1$ and $\rho_0 = 0$ are reference points corresponding to standard solution under full commitment and discretion, respectively. The full commitment solution features an initial interval of high but declining inflation ("start-up inflation") to temporarily sustain positive levels of output gap, which is desirable because a zero output gap is inefficiently low ($x^* > 0$). The discretionary solution is the well-known inflation bias.

Turning to the cases with interior values of initial reputation ρ_0, the most striking feature is that ρ_0 affects the optimal committed policy in a highly non-linear fashion. In particular, with high or moderate levels of initial reputation ($\rho_0 = .75,\ .5$), the start-up inflation is exaggerated relative to the full commitment case in the sense of a higher initial inflation and more aggressive declines of inflation over the first few periods, whereas with low levels of initial reputation ($\rho_0 = .2,\ .1$), the start-up inflation is mitigated or even reversed. Consequently, the deviation of output path from the full commitment case is small and controlled when the initial reputation is high or moderate, but is significant otherwise. The equilibrium reputation grows much slower with an initially high or moderate level than it does with a low initial level.

The key to understand the non-linear effect of ρ_0 is the endogenous response of the alternative inflation policy to the evolving state if the non-committed central banker is in place. Figure 2 plots the path of private sector's perceived alternative inflation policy and expected inflation associated with the equilibrium dynamics in Figure 1. First notice that the perceived alternative inflation policies with positive levels of initial reputation are all significantly below the inflation bias. With high levels of ρ_0, the initial alternative policy α_0 is half of the inflation bias, and with moderate or low levels of ρ_0, α_0 is three fourths of the inflation bias. This is because the private sector's expected inflation is anchored by the low or even negative committed inflation policy when they believe the central banker is committed with positive probability. The lower the expected inflation, the less tempting for the non-committed central banker to inflate the economy. With growing reputation, the perceived alternative policy quickly converges to half of the inflation bias.

Taking into account the effect of reputation on the alternative inflation policy, the committed central banker manages expected inflation to strike a balance between the low-inflation and high-output trade-off. To illustrate the trade-off, notice that the output is proportional to $a - e$, which in turn can be approximated by

$$\rho' (a - a') + (1 - \rho') (a - \alpha')$$

where the superscript prime is used to indicate the next-period variable. There are two channels that lead to a higher output: 1) a decline of committed policy; 2) a higher current-period committed policy than the next-period alternative policy.

When the initial level of reputation is high or moderate, a decline of committed policy is effective in stimulating the output because ρ' will also be high or moderate. The desirability of the second channel depends on whether α' can be low enough so that the level of a higher than α' does not cost too much welfare loss to the central banker. Recall that good reputation leads to a low α', hence the second channel is desirable when the initial reputation is high or moderate. Combining the first and second channels, it requires a higher a and a larger decline from a to a' to achieve the same output gain when the initial reputation is less than perfect but is good enough. Hence, the "start-up inflation" is exaggerated from the full commitment case.

When the initial level of reputation is low, the first channel will be ineffective and the second channel will be undesirable unless future reputation grows to be much better than the initial level. To grow reputation rapidly, it requires a large gap between a and α. With α being relatively high due to low levels of initial reputation, it is more efficient for the committed type to grow reputation with a level of a much lower than α. As a result, the start-up inflation is mitigated or even reversed, causing significant output loss in the first few periods after the replacement.

Across all cases with interior levels of initial reputation, the committed policy a gradually returns to the nearly-zero long-term rate after it hits a negative level and a further decline is no longer desirable. Along the returning path, the output gap is negative since $a - a' < 0$ and $a - \alpha' < 0$.

5.1.2. New regime with a non-committed central banker

The section above shows how the committed central banker's incentive to build reputation varies nonlinearly with his initial reputation. In this section, we study how the non-committed central banker spends his reputation along the transitional path after he takes the office.

Figure 3 plots the paths of the optimal *alternative* inflation policy (also the realized inflation rate), together with the implied output and reputation, in the first 4 years of a non-committed central banker's tenure. We plot the transitional dynamics conditional on the same 6 different levels of initial reputation as in Figure 1: 1('*'),.75('Δ'), .5('$\Diamond$'), .2('∇'), .1('+'), 0('$\bigcirc$').

The case with $\rho_0 = 0$ again corresponds to the solution under full discretion with the optimal alternative policy equal to the inflation bias. The case with $\rho_0 = 1$ is not consistent with the private sector's rational expectation but it serves as a limiting case where the reputation of the non-committed type is arbitrarily close to one. The optimal alternative policy in this case is around half of the inflation bias, thanks to a low expected inflation anchored by the perceived committed policy shown in Figure 4.

With interior levels of initial reputation, the transitional dynamics exhibits similar patterns across ρ_0. There is an initial interval of relatively low inflation at or below 3/4 of the inflation bias (1.5%), which is followed by a run-up of inflation towards the inflation bias. The higher the initial reputation is, the longer does the inflation stay low. When $\rho_0 = .75$, the period of low inflation lasts for three years. The initial interval of low inflation is associated with output boom whereas the run-up of inflation is associated with negative output gaps. The equilibrium reputation is, as expected, decreasing along the transitional path in all cases, yet the speed of decreasing is particularly slow if the initial reputation is good ($\rho_0 = .75$). It takes almost 4 years for the private sector to figure out the incumbent central banker is the non-committed type.

To understand this general pattern and especially why learning is so slow in the case of $\rho_0 = .75$, we plot in Figure 4 the private sector's perceived committed policy (it is also the announced policy) and expected inflation associated with the equilibrium dynamics in Figure 3. The perceived committed policy starts at the same level as in Figure 1 and is followed by an interval of increasingly aggressive disinflation policy. After the reputation declines to nearly zero, the perceived committed policy soars back to the inflation bias.

Notice that although the implied expected inflation is rising along the transitional path, its level is held lower than the realized inflation thanks to the interval of perceived disinflation policy. Believing that the incumbent central banker is likely to be the committed type leads the private sector to assign positive probability that the next-period inflation will be very low or negative. Even if there is observed discrepancy between the announced disinflation policy and the realized high rate in the current period, the private sector is willing to attribute the discrepancy more to the unexpected implementation error. When the discrepancy persists, the likelihood that it is caused by an i.i.d. error reduces and therefore the incumbent central banker's reputation declines. Recall from last section that when the reputation gets lower, the incentive for the committed central banker to grow reputation, if he were in place, starts to dominate his concern for output and therefore the optimal committed policy becomes lower or more negative. The larger discrepancy between the perceived committed policy and the realized inflation rate helps the private sector to learn faster and in turn the reputation declines faster. Lower levels of reputation makes it harder for the non-committed central banker to stimulate output so that the output gap also declines with the reputation. When the reputation becomes too poor, the committed central banker, if it were in place, would give up on growing it and hence implement a policy rate that is close to the inflation bias. When it is the case, the expected inflation jumps up to be above the realized inflation rate, leading to negative output gaps.

The particularly slow learning process in the case of $\rho_0 = .75$ is due to the prolonged interval of moderate disinflation policy that would be implemented by the committed central banker if he were in place. With good enough initial reputation, as explained in the last subsection, the committed central banker does not have strong incentive to implement a policy rate much lower than the alternative policy. As a result, having an inflation rate equal to the alternative policy does not cause much reduction in the central banker's reputation, and the committed central banker, if he were in place, would continue the moderate disinflation. So on and so forth until the central banker's reputation becomes low enough to incentivize the committed type to rapidly build reputation with more aggressive disinflation policies.

The equilibrium dynamics in the case of $\rho_0 = .75$, together with the perceived committed policy and the expected inflation, qualitatively replicates the inflation experience of the U.S. in 1960s and 1970s. In particular, the realized inflation started at a low level and increased continuously, which led to an economy boom in 1960s. The later stage of the inflation run-up was coupled by rounds of disinflation announcements made by the Federal Reserve bank, none of which were implemented. Eventually, the private sector's expected inflation soared and the economy ended with stagflation.

5.2. Dynamic response to shocks

In this subsection, we study the equilibrium response to two shocks in our model: the cost-push shock and the implementation error. The former is a supply shock whereas the latter can be interpreted as the demand shock as it moves the output and the inflation in the same direction. The two shocks also differ in their timing. The cost push shock is realized before the central banker takes policy action so there can be initial policy response to it. The demand shock is realized after the policy action so that all the policy response needs to come from expectation management.

Similar to the transitional dynamics, the equilibrium response to shocks is conditional on the type of central banker in place and his initial level of reputation. We assume that the realized implementation errors and cost-push shocks are zero in all periods except that in the initial period of the new regime, we let either the cost push shock or the demand shock takes place with a realized magnitude of one percent annually (0.25% quarterly). Recall that the cost-push shock has persistence .9 and the demand shock is transitory.

5.2.1. With a committed central banker

Figure 5 plots the impulse response (i.e. deviations from the transitional dynamics) to a positive persistent cost-push shock, conditional on the 6 different levels of initial reputation. In the full discretion solution ($\rho_0 = 0$), the inflation path simply reflects the persistence of the shock because the central banker has to split the shock's impact on inflation rate and output taking as given the expected inflation. In the full commitment solution ($\rho_0 = 1$), by contrast, the central banker responds to the shock with active expectation management. The inflation response takes the form of "flexible price-level targeting", namely the response is first positive and then negative, resulting in zero long-term effect of the cost-push shock on the price level. The active expectation management not only mitigates the initial inflation response to the shock, it also has output drop less on impact before it returns gradually to the normal level.

The inflation response with high or moderate levels of initial reputation ($\rho_0 = .75, .5$) is an exaggerated version of the full commitment solution in the sense that the initial inflation response is more accommodative but the following reduction in inflation is more aggressive.

As a result, the initial negative impact of the cost-push shock on output is similar to the full commitment case despite that the initial reputation is less than perfect. The output-driven policy response echoes the central banker's strong incentive to stimulate output in the transitional dynamics when he starts with high or moderate levels of initial reputation. Similarly, when the output is the main concern, the reputation is the by-product of the inflation response. In this case, there is significant reputation gain from the aggressive decline of inflation following the initial accommodation.

When the central banker starts with low levels of initial reputation ($\rho_0 = .2, .1$), interestingly, the initial policy response turns out to be even more accommodating than the full discretion solution. Moreover, it is followed by aggressive disinflation until the inflation reaches the similar level as the full commitment solution. Such a path of inflation greatly mitigates the negative effect of the cost-push shock on output. The initial output loss in both cases of ρ_0 is smaller than the full commitment case. In the case of $\rho_0 = .2$, the entire path of output lies above the full commitment case. The mitigated output loss does not come with no cost. The reputation in both cases of ρ_0 exhibits downward deviation from its transitional path, indicating a slower learning by the private sector. Recall that, in absence of shocks, the committed central banker with low levels of initial reputation incurs significant output loss in order to grow reputation rapidly. When the cost-push shock occurs, a further drop of output becomes increasingly undesirable for the committed central banker. Hence, his policy response is to mitigate further output loss with a slower growth of reputation.

Figure 6 plots the impulse response to a positive transitory demand shock with the same 6 different levels of initial reputation. A demand shock is unexpected so there cannot be any contemporaneous policy response. A positive demand shock, given the expected inflation, stimulates output. To smooth the effect on output over time, the only channel is through managing expected inflation, namely by promising a higher-than-avearge but declining inflation. The cases with $\rho_0 = 1$ and $\rho_0 = 0$ illustrate the two extreme cases in terms of the central banker's leverage over expected inflation. In the full discretion solution ($\rho_0 = 0$), there is no inflation response or output smoothing since the central banker has no leverage over expected inflation. In the full commitment solution ($\rho_0 = 1$), the optimal policy mitigates the initial impact of the shock on output and makes the effects of the transitory shock persist, generating an interval of output boom similar to the effect of "start-up inflation".

Now we turn to the cases with interior levels of initial reputation. First, the initial impact of demand shock on the reputation is pre-determined. With high or moderate levels of initial reputation ($\rho_0 = .75, .5$), the impact is positive because in the transitional dynamics, the committed inflation policy is higher than the perceived alternative one in the initial period. Hence, a positive surprise in realized inflation speeds up the private sector's learning. With low levels of initial reputation ($\rho_0 = .2, .1$), the demand shock's impact on reputation is significantly negative because in the transitional dynamics, the committed central banker rapidly builds his reputation with a policy much lower than the perceived alternative one. Second, taking into account the initial impact of demand shock on the reputation ρ_1, the policy response replicates the nonlinear effect of reputation on committed policy in the transitional dynamics. In particular, with high or moderate levels of reputation ρ_1, the committed central banker mainly concerns about output. As a result, his policy response is an exaggerated version of the full commitment solution with a higher period-1 inflation followed by more aggressive declines of inflation, to make up for his less-than-perfect control over expected inflation. The implied output path largely replicates the full commitment case. With low levels of reputation ρ_1, the committed central banker mainly concerns the reputation growth so his policy response is to make up for the reputation loss caused by the demand shock. Consequently, there is a prolonged interval of downward deviation of output from its transitional dynamics.

5.2.2. With a non-committed central banker

When the central banker is the non-committed type, the equilibrium policy responses to shocks are more passive than they are with a committed central banker. This is because the non-committed central banker does not have strategic power over reputation evolution so the optimal alternative policy is determined by the perceived committed policy response to the shock and its implied reputation change from the transitional path. With interior levels of initial reputation, the first order effect of the shock is through the implied reputation change. This is because the transitional dynamics shares the same pattern across various levels of initial reputation and differs only in how fast the reputation depletion triggers the run-up of inflation by the non-committed central banker and in turn the length of the output boom. Any change in reputation due to the shock is going to delay or speed up the reputation depletion and therefore affect the global dynamics of equilibrium variables. To illustrate this global effect of shock, we plot the equilibrium dynamics with the shock rather than its deviation from the transitional dynamics.

Figure 7 plots the equilibrium dynamics with a positive persistent cost-push shock taking place in the initial period. The case with $\rho_0 = 0$ is the same as in Figure 5, where the inflation response simply reflects the persistence of the shock. Interestingly, the case with $\rho_0 = 1$ shares the same pattern with the full discretion solution except that the magnitude of inflation response is smaller, thanks to the anchoring effect of the perceived committed policy (as in Figure 5) on the expected inflation. With high and moderate levels of initial reputation ($\rho_0 = .75, .5$), the private sector's learning is faster because the perceived policy response by the committed central banker (similar to that in Figure 5) should be a large decline of inflation after the initial accommodation of the shock, whereas the realized inflation only responds with a moderate decline because the non-committed central banker does not honor the promised large decline of inflation. With low levels of initial reputation ($\rho_0 = .2, .1$), the private learning is slower because the perceived committed policy is very accommodating to the shock on impact. By contrast, the alternative policy needs not to be very accommodating since there is an output boom in absence of the shock. Recall that in transitional dynamics, the committed policy is much lower than the alternative policy so accommodating the cost-push shock makes the two policies closer and thus the realized inflation less revealing about the central banker's type.

Turning to the positive transitory demand shock that takes place in the initial period, Figure 8 plots the equilibrium dynamics. Notice that in the cases with $\rho_0 = 1$ and $\rho_0 = 0$, there is no response of the equilibrium inflation policy to the shock because the non-committed central banker does not have strategic power over the expected inflation. When the initial reputation takes interior values, the demand shock's initial impact on reputation is the same as in Figure 6. With high or moderate levels of initial reputation ($\rho_0 = .75$, $.5$), the impact is positive. With low levels of initial reputation ($\rho_0 = .2$, $.1$), the impact is negative. The initial positive impact on reputation prolongs the interval of low inflation policy and hence slows down the private sector's learning along the entire transitional path. The initial negative impact on reputation speeds up the convergence of inflation policy to the inflation bias and in turn makes the private sector's learning faster along the transitional path.

6. Conclusions and Final Remarks

In this paper, we study optimal monetary policy in setting where the private sector is forwarding-looking and learning about the type of central bank in place. The central bank can be either a patient and committed type, or an impatient type that cannot commit. Being able to commit or not, the central bank in place chooses inflation policies optimally, taking into account the learning and rational expectation of the private sector. To solve the model, we adopt a mechanism design approach so that the equilibrium of the model can be obtained as the solution to a recursive optimization problem of the committed central bank. In particular, we view the committed type as choosing its action and the action of the non-committed type subject to an incentive compatibility constraint. The relevant incentive compatibility constraint is an unusual one because it involves the beliefs of a third party, i.e., the private sector.

The numerical solution to a calibrated model reveals two sets of results that capture various aspects of the historical behavior of monetary policy, inflation and real activity. If the central bank in place is the non-committed type yet the private sector attaches a high likelihood to there being a committed type in place, there will be lengthy intervals of high real activity with gradually rising actual and expected inflation. As the history of positive inflation surprises leads to a decline in the likelihood that a committed type is in place, the non-committed central bank will optimally choose to experience a rapidly rising inflation accompanied by declining real activity (stagflation). Ultimately, its type will be revealed and it will simply experience high and volatile inflation with little effect on real activity. If the central bank in place is the committed type, its optimal monetary policy will be similar to the standard solution under full commitment, conditional on the private sector attaching a high likelihood to there being a committed type in place. In this case, there will be "great moderation" with declining inflation and stable real activity. But if the private sector attaches a low likelihood to there being a committed type in place, the committed type will optimally adopt anti-inflation policies with real costs to build its reputation rapidly. The response of inflation to the energy price shock will also be time-varying (in particular, the response becomes smaller) over the duration of the regime as the private sector gradually learns that it is the committed type in place.

In the current paper, we assume that the central bank can directly control inflation (up to an implementation error). Although it is a simplifying modelling choice, we can work out implications for interest rates that come directly out of this setup. For example, introducing an IS curve and the Fisher equation would allow us to determine the short-term nominal interest rate from the output and the inflation target; and the private sector's learning would allow us to determine the long-term nominal interest rate as it is dominated by long-term expectations of inflation. Deriving the relations between the short-term and long-term interest rates can be a valuable future extension of the current work.

Another strong assumption in our current model is that the non-committed policy authority is myopic. To allow for a patient non-committed type, one needs to modify the recursive optimization of the committed type to incorporate the promised value for the non-committed type in the incentive compatibility constraint. This will significantly complicate the computation. We leave this extension to future work.

References

Atkeson, A. and Kehoe, P. J. (2006), 'The advantage of transparency in monetary policy instruments', *Federal Reserve Bank of Minneapolis Staff Report 297* .

Backus, D. A. and Driffill, J. (1985*a*), 'Inflation and Reputation.', *American Economic Review* **75(3)**, 530–538.

Backus, D. and Driffill, J. (1985*b*), 'Rational expectations and policy credibility following a change in regime', *The Review of Economic Studies* **52**(2), 211–221.

Barro, R. J. (1986), 'Reputation in a model of monetary policy with incomplete information', *Journal of Monetary Economics* **17**(1), 3–20.

Chang, R. (1998), 'Credible monetary policy in an infinite horizon model: Recursive approaches', *Journal of Economic Theory* **81**, 431 – 461.

Cukierman, A. and Liviatan, N. (1991), 'Optimal accommodation by strong policymakers under incomplete information', *Journal of Monetary Economics* **27**(1), 99–127.

Cukierman, A. and Meltzer, A. H. (1986), 'A theory of ambiguity, credibility, and inflation under discretion and asymmetric information', *Econometrica* **54**(5), 1099.

Debortoli, D. E. and Nunes, R. (2014), 'Monetary regime switches and central bank preferences', *Journal of Money, Credit and Banking* **46**(8), 1591–1626.

Eusepi, S. and Preston, B. (2018), 'The science of monetary policy: An imperfect knowledge perspective', *Journal of Economic Literature* **56**(1), 3–59.

Evans, G. W. and Honkapohja, S. (2003), 'Adaptive learning and monetary policy design', *Journal of Money, Credit and Banking* **35**(6), 1045–1072.

Evans, G. W. and Honkapohja, S. (2008), Expectations, Learning, And Monetary Policy: An Overview Of Recent Research, Working Papers Central Bank of Chile 501, Central Bank of Chile.

Faust, J. and Svensson, L. E. O. (2001), 'Transparency and credibility: Monetary policy with unobservable goals', *International Economic Review* **42**(2), 369–397.

Gali, J. (2015), *Monetary Policy, Inflation, and the Business Cycle*, Princeton Univers. Press.

Gali, J. and Gertler, M. (1999), 'Inflation dynamics: A structural econometric analysis', *Journal of Monetary Economics* **44**(2), 195–222.

Golosov, M., Tsyvinski, A. and Werquin, N. (2016), Recursive contracts and endogenously incomplete markets, *in* 'Handbook of Macroeconomics', Elsevier, pp. 725–841.

Hansen, S. and McMahon, M. (2016), 'First impressions matter: Signalling as a source of policy dynamics', *The Review of Economic Studies* **83**(4), 1645–1672.

Khan, A., King, R. G. and Wolman, A. L. (2003), 'Optimal monetary policy', *Review of Economic Studies* **70**(4), 825–860.

King, R. G., Lu, Y. K. and Pastén, E. S. (2008), 'Managing expectations', *Journal of Money, Credit and Banking* **40**(8), 1625–1666.

Kydland, F. E. and Prescott, E. C. (1980), 'Dynamic optimal taxation, rational expectations and optimal control', *Journal of Economic Dynamics and Control* **2**, 79–91.

Lu, Y. K. (2013), 'Optimal policy with credibility concerns', *Journal of Economic Theory* **148**(5), 2007–2032.

Lu, Y. K., King, R. G. and Pasten, E. (2016), 'Optimal reputation building in the new keynesian model', *Journal of Monetary Economics* **84**, 233–249.

Mailath, G. J. and Samuelson, L. (2006), *Repeated Games and Reputations: Long-run Relationships*, Oxford University Press.

Marcet, A. and Marimon, R. (2017), 'Recursive contracts', *Working paper* .

Matthes, C. (2015), 'Figuring out the fed-beliefs about policymakers and gains from transparency', *Journal of Money, Credit and Banking* **47**(1), 1–29.

Mishkin, F. S. and Schmidt-Hebbel, K. (2007), Does Inflation Targeting Make a Difference?, NBER Working Papers 12876, National Bureau of Economic Research, Inc.

Osborne, M. J. and Rubinstein, A. (1994), *A Course in Game Theory*, MIT Press Ltd.

Phelan, C. and Stacchetti, E. (2001), 'Sequential equilibria in a ramsey tax model', *Econometrica* **69**(6), 1491–1518.

Roger, S. and Stone, M. R. (2005), 'On target? the international experience with achieving inflation targets', *IMF Working Paper No. 05/163* .

Schaumburg, E. and Tambalotti, A. (2007), 'An investigation of the gains from commitment in monetary policy', *Journal of Monetary Economics* **54**(2), 302–324.

Walsh, C. E. (2010), *Monetary Theory and Policy*, MIT Press Ltd.

Woodford, M. (2003), 'Optimal interest-rate smoothing', *The Review of Economic Studies* **70**(4), 861–886.

Woodford, M. (2013), 'Macroeconomic analysis without the rational expectations hypothesis', *Annual Review of Economics* **5**(1), 303–346.

Xandri, J. P. (2017), 'Credible reforms: A robust implementation approach', *Working paper* .

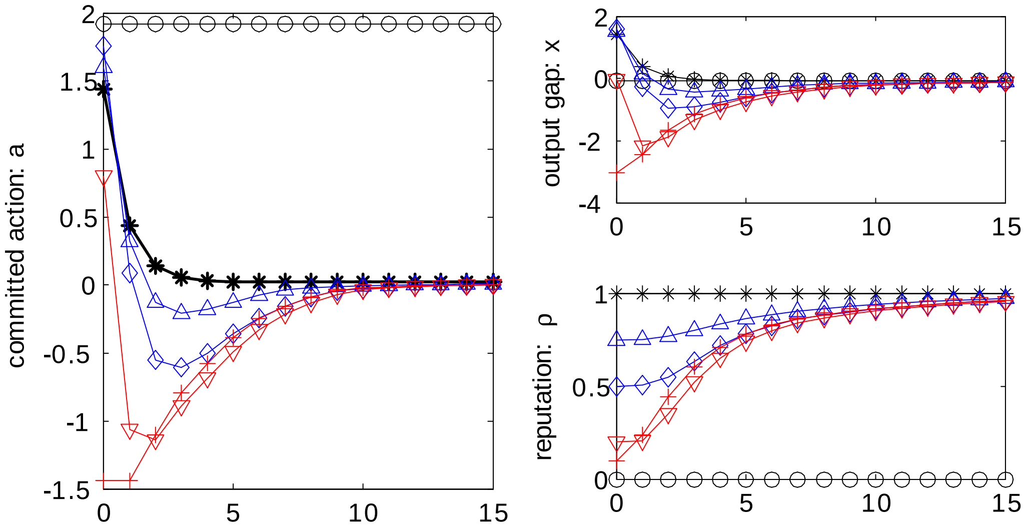

Figure 1: Transitional dynamics: new regime with committed central banker. Each path is plotted based on one of the six levels of initial reputation: 1('*'), .75('Δ'), .5('$\Diamond$'), .2('∇'), .1('+'), 0('$\bigcirc$'). Committed policy action is annualized mean inflation rate. Output gap is in percent deviation from its distorted steady state. Reputation is the private sector's belief that the current central banker is the committed type.

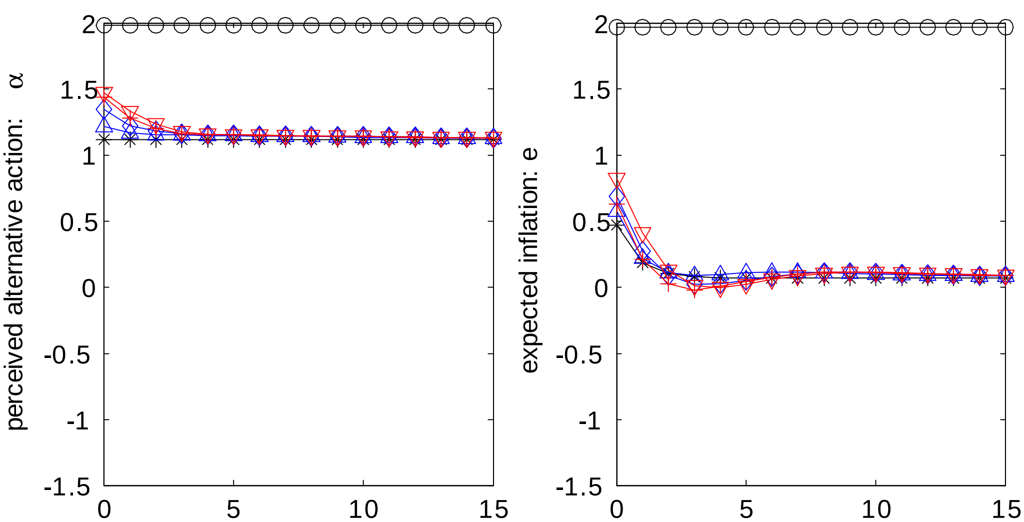

Figure 2: Transitional dynamics: new regime with committed central banker. Each path is plotted based on one of the six levels of initial reputation: 1('*'), .75('Δ'), .5('$\Diamond$'), .2('∇'), .1('+'), 0('$\bigcirc$'). Perceived alternative policy action is the annualized mean inflation rate if the current central banker were the alternative type. Expected inflation is the private sector's expected next-period inflation in annualized rate.

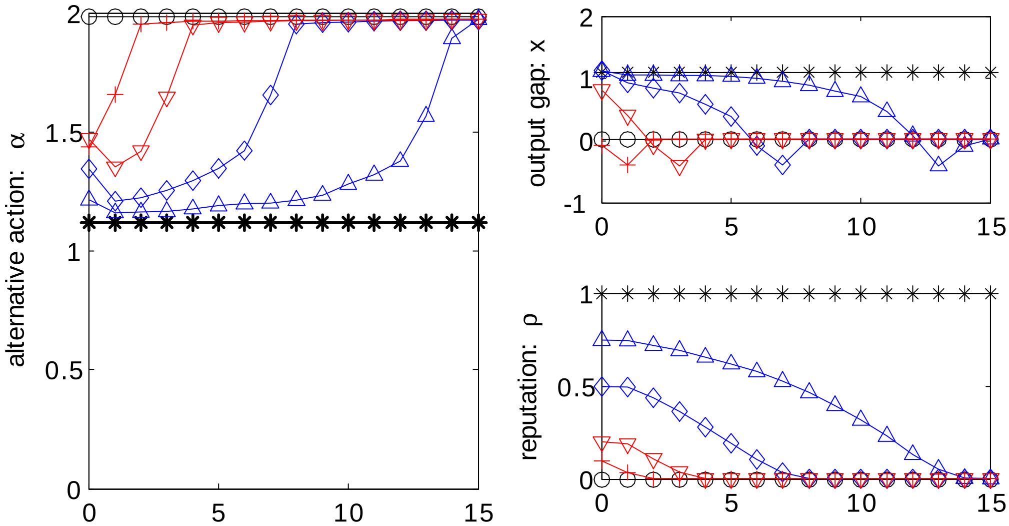

Figure 3: Transitional dynamics: new regime with non-committed central banker. Each path is plotted based on one of the six levels of initial reputation: 1('*'), .75('Δ'), .5('◊'), .2('∇'), .1('+'), 0('◯'). Alternative policy action is annualized mean inflation rate. Output gap is in percent deviation from its distorted steady state. Reputation is the private sector's belief that the current central banker is the committed type.

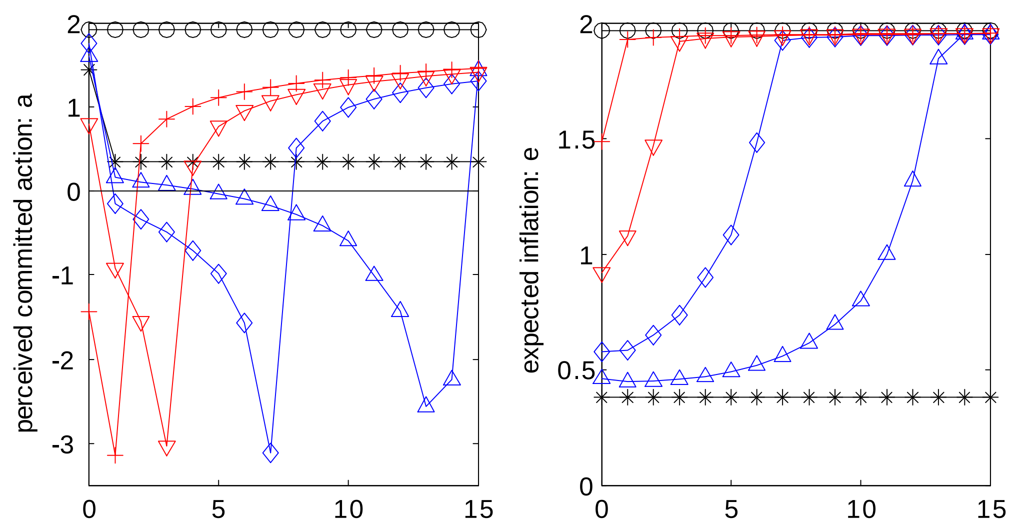

Figure 4: Transitional dynamics: new regime with non-committed central banker. Each path is plotted based on one of the six levels of initial reputation: 1('*'), .75('Δ'), .5('◊'), .2('∇'), .1('+'), 0('◯'). Perceived committed policy action is the annualized mean inflation rate if the current central banker were the committed type. Expected inflation is the private sector's expected next-period inflation in annualized rate.

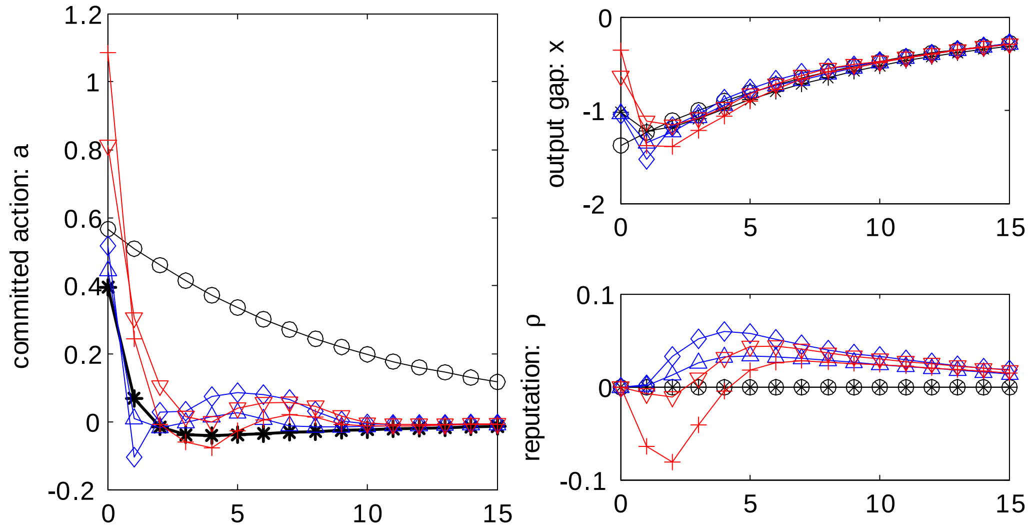

Figure 5: Impulse response with a committed central banker to a persistent (0.9) cost-push shock (one percent annually) All variables are plotted as deviations from the transitional dynamics. Each path is plotted based on one of the six levels of initial reputation: 1('*'), .75('Δ'), .5('$\Diamond$'), .2('∇'), .1('+'), 0('$\bigcirc$').

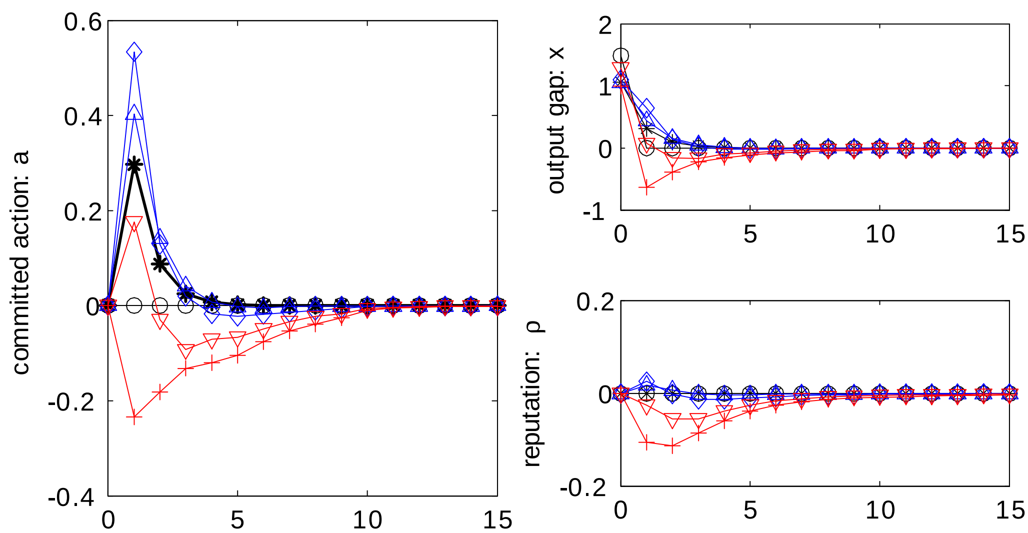

Figure 6: Impulse response with a committed central banker to a one-time demand shock (one percent annually) All variables are plotted as deviations from the transitional dynamics. Each path is plotted based on one of the six levels of initial reputation: 1('*'), .75('Δ'), .5('$\Diamond$'), .2('∇'), .1('+'), 0('$\bigcirc$').

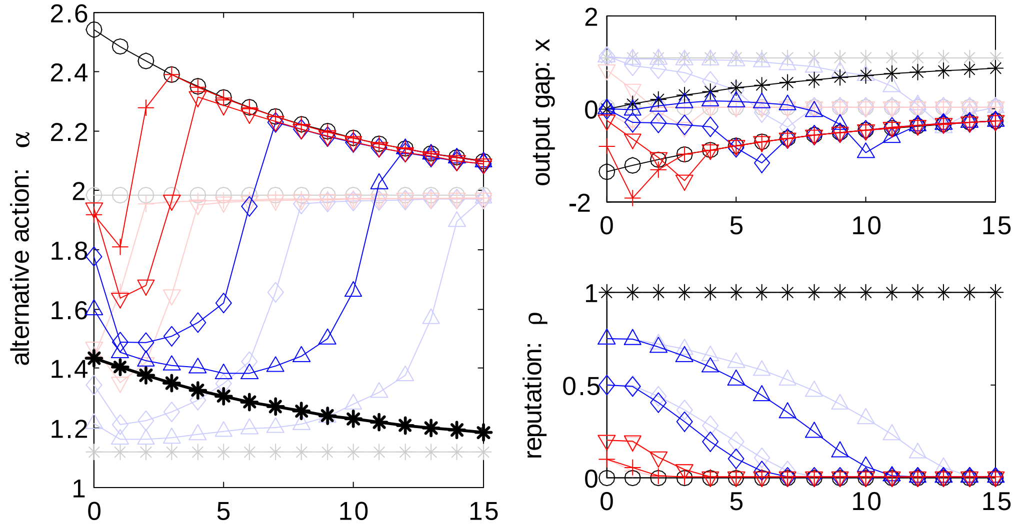

Figure 7: Equilibrium dynamics with a non-committed central banker in the presence of a persistent (0.9) cost-push shock (one percent annually) in period 0. The shaded path is the transitional dynamics without the shock. Each path is plotted based on one of the six levels of initial reputation: 1('*'), .75('Δ''), .5('$\Diamond$'), .2('∇'), .1('+'), 0('$\bigcirc$').

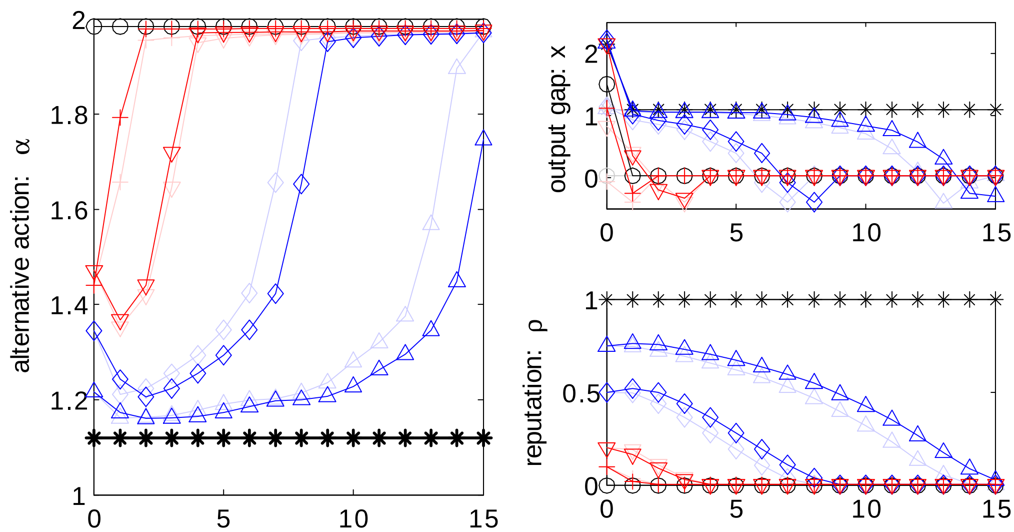

Figure 8: Equilibrium dynamics with a non-committed central banker in the presence of a one-time demand shock (one percent annually) in period 0. The shaded path is the transitional dynamics without the shock. Each path is plotted based on one of the six levels of initial reputation: 1('*'), .75('Δ'), .5('$\Diamond$'), .2('∇'), .1('+'), 0('$\bigcirc$').