Policy design with private sector skepticism in the textbook New Keynesian model

Robert G. King   Yang K. Lu   Ernesto S. Pasten

May 2012

Abstract

How should policy be optimally designed when a monetary authority faces a private sector that is somewhat skeptical about policy announcements and which interprets economic data as providing evidence about the monetary authority's preferences or its ability to carry through on policy plans? To provide an answer to this question, we extend the standard New Keynesian macroeconomic model to include imperfect inflation control (impementation error relative to an inflation action) and Bayesian learning by private agents about whether the monetary authority is strong (capable of following through on announced plans) or weak (producing higher and more volatile inflation). In a benchmark case, we find that optimal policy involves dramatic anti-inflation actions which include an interval of deflation during the early stages of a plan, motivated by investing in a reputation for strength. Such policies resemble recommendations during the 1980s for a "cold turkey" approach to disinflation. However, we also find that such policy is not robustly optimal. A more "gradualist" policy arises if the initial level of credibility is very low. We also investigate a setting where the alternative (weak) monetary authority follows a simple behavioral rule that mimics variations in the strong authority's policy action but with a bias toward higher and more volatile inflation. In this case, which we call a "tag along" alternative policymaker, a form of gradualism is always optimal. This finding suggests the importance of extending the analysis of optimal policy design to a setting with strategic interaction between types of policymakers, which is the focus of our ongoing work.

1 Introduction

Policy design in modern dynamic stochastic general equilibrium models with nominal frictions is conducted in one of two modes: the monetary authority is fully capable of commitment or completely unable to commit. The implications of optimal monetary policy for the average level of inflation and the response of inflation to shocks are substantially different under these two assumptions. The commitment solution, in particular, features detailed dynamic plans for the evolution of inflation over time which have a flexible price level targeting interpretation. Since most monetary authorities believe that their policy plans will be carried out, they are led toward some version of the commitment solution as a guide to the design of policy. By contrast, in an equilibrium without commitment, there are alternative simpler rules which give rise different macroeconomic outcomes, with higher average inflation and permanent variations in the price level.

But how should policy be designed in the middle ground where monetary authorities frequently find themselves, which is that they face a private sector that is somewhat skeptical about policy announcements and which interprets economic data as providing evidence about the monetary authority's preferences or its ability to carry through on policy plans?

In this paper, we provide a reference answer to this question, studying a version of the textbook New Keynesian monetary policy model of linear quadratic form commonly used to simply represent the richer macroeconomic dynamics of medium scale policy models. Within the well-known full commitment case, in which the detailed dynamic policy plan is fully credible, this model has two striking features[1] First, optimal policy involves an initial interval of high, but declining inflation that stimulates real activity, which we term the "start up" phenomenon. Second, optimal policy involves significant accommodation of inflation shocks, so as to offset consequences for real economic activity. A by-product of our work with this model involves showing that these core implications carry over to a setting in which the monetary authority can only imperfectly control inflation due to an implementation error.

To answer the question posed above, it is necessary for us to be specific about the natures of the monetary authority in place and the alternative which the private sector believes may be present. In doing so, we draw a distinction between a strong

[1] See, for example Clarida, Gali, Gertler [1999] and King and Wolman [1999]

monetary authority – one that is capable of formulating and carrying out the type of detailed plan derived in a commitment equilibrium – and the extent of its credibility. Our initial focus is on deriving the optimal policy for an authority that can commit but faces a private sector that attaches a probability – an extent of credibility – to that the announced policy plan but also believes that policy may be selected according to an alternative plan that is more inflationary. In terms of the behavior of the alternative monetary authority, our reference case is that it follows the simple rule arising in the full information equilibrium without policy commitment, which is well understood to involve both inflation and stabilization biases but no start-up period. When the credibility index is constant over time, we find that optimal policy for the strong monetary authority involves mixture of standard commitment and non-commitment outcomes: there is a degree of inflation bias and stabilization bias, in addition to a start-up period of scale diminishing in credibility. This fixed credibility case is a second by-product of our work and, like the textbook linear-quadratic model, there is a closed form solution for optimal policy and macroeconomic outcomes.

With this background, we then construct our main environment in which the strong authority formulates an optimal dynamic policy plan but does so recognizing that (i) actual inflation outcomes are more variable than its policy choices due to implementation error; and (ii) private agents learn from inflation outcomes about the nature of the authority that is in place. Under our assumption that the monetary authority has full commitment capability, but that price setters and other private agents are skeptical, we find important departures from standard conclusions about startup inflation and responses to inflation shocks. First, from relatively high levels of credibility (50% and higher), optimal policy features an initial interval of lower inflation than the full commitment solution, with the nature of the path depending on the initial extent of private sector skepticism but frequently involving deflation. Private sector learning is not immediate due to the presence of implementation error, which masks the policy actions of the monetary authority. Essentially, the monetary authority engages in an initial period of *reputation building*.[2] Second, optimal policy

[2]In the sense of Cripps, Mailath and Samuelson [2004] and Lu [2012]. The precise implications of this reputation investment for optimal policy depend on the structure of the economy, including the learning rule of the private sector, so that there is not a simple, comprehensive prescription such as the use of the "timeless perspective" advocated by Woodford [1999]. Kurozumi [2008] and Loisel [2008] study the issue of whether optimal monetary policy is sustainable in the sense of Chari and Kehoe [1990], using a different notion of reputational equilibria.

features a time-varying response to inflation shocks. Generally, in our basic model, learning is very fast if initial reputation/credibility is relatively high and so the monetary authority's optimal policies typically approach the standard full commitment solution fairly quickly.[3] Hence, with endogenous credibility and from relatively high initial levels of credibility, the short-run is dominated by reputation investment and the long-run looks like the flexible price level targeting solution familiar from the literature.

In the late 1970s and early1980s, there was much discussion of the appropriate strategy for disinflation in the United States and in other countries. One approach was *gradualism*, by which policy reduced the inflation rate slowly with the objective of producing small real losses. Another was to undertake a rapid disinflation, frequently called the *cold turkey* strategy.[4] A newly reorganized monetary authority with full commitment and credibility in the New Keynesian model would adopt a gradualist policy and there would be a resulting boom in real economic activity.[5] By contrast, with endogenous credibility and starting from relatively high levels, our reference analysis shows that a rapid disinflation is optimal and that there is a recession, whose depth and duration is larger when initial credibility is lower. Further, our reference analysis shows that policy responds aggressively to avoid deterioration of credibility in the face of implementation errors and other price shocks, with the intensity again depending on the level of credibility. However, if we consider a monetary authority which starts with a relatively low level of credibility (25%), then the cold turkey strategy is not optimal. Instead, the authority behaves in a gradualist manner, reducing inflation while balancing the real costs of disinflation with the gains from investing in reputation. Even though it is gradualist, optimal policy does bring about a deep recession. Further, positive implementation errors – inflation high relative to the optimal action – lead to a more protracted interval of gradual disinflation with higher real costs.

Moreover, an important literature in the 1980s on credible control of inflation stressed that monetary authorities not capable of full commitment (colloquially, weak)

[3] The deflationary interval and the rapid learning are results which are also obtained by Cogley, Matthes and Sbordonne [2011] in substantively related research using a different computational approach.

[4] These two strategies are discussed, for example, by Sargent [1982] and Bernanke [2004].

[5] Ball [1994]. Note that this phenomenon extends to our constant partial credibility extension in section 3 below.

might be induced to deliver low and stable inflation by the force of trigger-strategy expectations[6]. In essence, a weak monetary authorities would find it desirable to *mimic* the behavior of a strong monetary authority – specifically, one that was modeled as mechanically a low inflation policy – due to the threat that it would be permanently faced with high inflation expectations and have to repeatedly accommodate these.

As an initial exploration of the potential consequences of such mimicking, we study the nature of optimal policy when the alternative weak monetary authority takes a policy action that adds a "time-varying inflation premium" to the strong central bank's optimal policy action. This *tag-along* behavior is known to the private sector, so that it affects the learning rule facing a strong monetary authority in its optimal policy design. This modification has a dramatic effect on optimal inflation policy: it is gradualist from all levels of initial credibility, although the motivations for the measured pace differ. At a high initial level, the optimal policy closely resembles the commitment solution – learning plays little role – and there is positive but declining inflation with an initial interval of real stimulus. For a low level, the optimal policy also involves an initial interval of high, but declining inflation, but for a very different reason: the monetary authority correctly understands that private sector inflation expectations will be stubborn in response to disinflation events. There is a u-shaped recession in real economic activity.

To generate the results in the paper, we draw on a recursive characterization of optimal policy problems with a forward-looking expected value constraints that we derive in companion work[7]. Appropriate for a setting of imperfect public monitoring, this recursive problem features a credibility/reputation state variable, as well as various commitment multipliers, and is applicable to settings in which the alternative weak decision-maker optimizes on a period-by-period basis as well as the basic cases that we explore here. It builds on the basic approach to optimal policy design with forward-looking constraints developed in work on dynamic contracts[8] and, more specifically, on the recursive approach to monetary policy design.[9]

The organization of the paper is as follows. In section 2, we describe our variant of the textbook New Keynesian model and lay out the recursive optimal policy problem..

[6]Sometimes also described as reputation.

[7]King, Lu, and Pasten [2011].

[8]By the literature that follows the path laid out by Marcet and Marimon [1998, 2011]

[9]See Khan, King and Wolman [2003] for an early analysis of a New Keynesian model along recursive lines.

In section 3, we study the case in which credibility is fixed. In section 4, we analyze the reference case in which the alternative policymaker is an automaton, by which we mean that it follows the simple policy rule that would be optimal behavior in an equilibrium without commitment. In section 5, we study the tag-along case. In section 6, we provide a summary, conclusions, and overview of planned future work.

2 Recursive optimal policy

A standard New Keynesian optimal policy problem involves a monetary authority maximizing an expected present discounted value objective such as

$$\max_{\{\pi_t, x_t\}_{t=0}^{\infty}} E_0\{\sum_{t=0}^{\infty} \beta^t u(\pi_t, x_t)\} \tag{1}$$

defined over inflation π and output x (relative to an efficient level x^*). Typically, the momentary objective is assumed to be quadratic, as in

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

with $h > 0$. Output is a good and inflation is a bad at small positive values of x and π, in the sense that $u_\pi = -\pi < 0$ and $u_x = -h(x - x^*) > 0$.

The standard NK constraint is a forward-looking specification for inflation,

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

for each period $t = 0, 1, ...\infty$. In this expression, as is also standard, we include a shock to inflation ς_t governed by an exogenous Markov process.[10]

[10] Note that these specification have four properties that are central to understanding the dynamic behavior of optimal policy. First, as stressed by Ball [**1994**], a perfectly credible anticipated disinflation raises output (directly from (3)). Second, output in the economy is inefficiently low (there are losses $h(x_t - x^*)^2$ in the momentary objective (2)). The combination of these two properties means that it is desirable to have an initial interval of high, but declining inflation, as part of an optimal policy plan. Third, as stressed by Ball [1995], an imperfectly credible disinflation can readily yield a contraction in output (an implication of (3)). For example, Goodfriend and King [2005] show that a gradual decline in inflation coupled with expectations that inflation will remain at a high initial level will lead to an intensifying recession. Fourth, there are costs to both inflation and output deviations in (2)). This last feature governs the efficient extent of initial "start up" inflation. It also circumscribes the response of the economy to inflation shocks (ς_t), implying that it it is not desirable to fully stabilize either output or inflation

We present the elements of this familiar model in a relatively terse manner, so that Table 2.1 provides the reader with a list of notation and definitions.

2.1 Modifications of familiar model

Relative to the literature, we introduce three complications to this basic model.

2.1.1 Implementation error in inflation policies

At the start of each period, a policy action a_t must be taken on the basis of past inflation and the current inflation shock ς_t. Period t inflation is then generated stochastically according to

$$\pi_t = a_{\tau t} + \varepsilon_t \tag{4}$$

where ε_t is an implementation error with a zero mean random variable and a finite variance.[11] The action a depends on the type of the monetary authority, τ, but the distribution of the implementation error does not.

2.1.2 A two-stage economic structure

We assume that inflation expectations are formed and output is determined after the inflation outcome is realized, so that our introduction of implementation error is associated with a two-stage structure that mandates breaking each date t into subperiods.

The monetary authority announces its planned action at the start of the period, after viewing the shock ς_t and the realization of inflation from the prior period π_{t-1}.

2.1.3 Partial credibility

Through the paper, we view private agents as forming expected inflation after seeing π_t and with a degree of skepticism about whether inflation will be generated according to the monetary authority's announced plan a or otherwise. That is, the private sector is uncertain about which of two monetary authority types is managing the economy. The strong type $\tau = 1$ is capable of commitment: it makes an optimal choice on policy plans in period zero and commits to the plans in all subsequent periods; our

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

objective to characterize these optimal plans. The alternative type $\tau = 2$, however, uses an exogenous rule in choosing its policy actions, which will be specified further below.

In terms of the dynamics of *reputation/credibility*, we study two cases that are discussed next.

Partial credibility without learning In section 3, we assume with a fixed level of credibility ρ for the monetary authority's plan. That is, we explore

$$\pi_t - \kappa x_t - \beta[\rho E_t \pi_{t+1} | a_1(s_{t+1}) + (1-\rho) E_t \pi_{t+1} | a_2(s_{t+1})] - \varsigma_t \geq 0 \qquad (5)$$

for each period $t = 0, 1, ...\infty$. In this expression, ρ is the private sector's fixed probability that future inflation will be generated by the actions of the strong type (type $\tau = 1$) according to its optimal plan a_1, which depends on the as yet unspecified future state of the economy s_{t+1} which includes π_t. Thus, $(1-\rho)$ is the complementary probability that future inflation will be generated by the actions of the alternative type (type $\tau = 2$) according to an exogenous rule a_2, which also can depend on the future state of the economy s_{t+1}.[12] This initial specification has the convenient property that the full commitment solution arises as a special case when $\rho = 1$, the no commitment solution occurs when $\rho = 0$ and consideration of intermediate cases is also straightforward.

Partial credibility with learning In sections 4 and 5, we assume that there is Bayesian learning about the monetary authority's type. When the current inflation rate is observed, the private sector's probability (as of the start of period t) that the monetary authority is of type 1 is updated according to a Bayesian learning specification

$$\rho_{t+1} = b(\pi_t, \rho_t; a_1, a_2) \qquad (6)$$

where the conventional form of the Bayesian updating function b will be detailed below.

[12] The specification is set up so that we can endogenize the strategy of the alternative type, utilizing the approach developed in KLP [2011]. However, it is beyond the scope of the current paper.

The forward-looking specification for inflation is then

$$\pi_t - \kappa x_t - \beta[\rho_{t+1} E_t \pi_{t+1} | a_1(s_{t+1}) + (1-\rho_{t+1}) E_t \pi_{t+1} | a_2(s_{t+1})] - \varsigma_t \geq 0 \quad (7)$$

again for each period $t = 0, 1, ...\infty$. In this expression, ρ_{t+1} is the private sector's updated probability that the monetary authority is of type 1 and $1 - \rho_{t+1}$ is the updated probability attached to an alternative type 2.

2.2 The alternative type and expected inflation

In this paper, we adopt two specifications for the alternative type.

One is a traditional specification: we assume that the alternative type is an *automaton*, following a simple rule, which we write as $a_2 = \mu + \phi\varsigma_t$. However, we choose the parameters μ and ϕ so that the type $\tau = 2$ actions correspond to those that would be optimal if there were no commitment.

The other is more unusual: we assume that the alternative type adds a "time-varying inflation premium" to the strong type's policy action: $a_2 = a_1 + \mu + \phi\varsigma_t$. We call this a *tag-along* policymaker.

The expected inflation implications of both of these specifications can be captured within the general form $a_2 = \omega a_1 + \mu + \phi\varsigma_t$. Given the behavior of the alternative type, the private sector's expected inflation then is:

$$\begin{aligned} E_t\pi_{t+1} &= \rho_{t+1}[E_t\pi_{t+1}|a_1(s_{t+1})] + (1-\rho_{t+1})[\omega(E_t\pi_{t+1}|a_1(s_{t+1})) + \mu + \phi E_t\varsigma_{t+1}] \\ &= l_{t+1}[E_t\pi_{t+1}|a_1(s_{t+1})] + (1-\rho_{t+1})[\mu + \phi E_t\varsigma_{t+1}] \end{aligned}$$

with $l = \rho + (1-\rho)\,\omega$ defined for convenience. Note that $E_t\pi_{t+1}|a_1(s_{t+1}) = E_t a_1(s_{t+1})$, given that there is zero expected implementation error, so that l_{t+1} captures the degree of control that the monetary authority has over near-term expected future inflation, which we colloquially refer to as its *leverage* on expectations. Note also that a part of near-term expected future inflation $(1-\rho_{t+1})[\mu + \phi E_t\varsigma_{t+1}]$ is beyond the control of the monetary authority: this exogenous component is larger if policy is less credible (lower ρ) and if the autonomous component of near-term expected inflation is larger.

2.3 Recursive optimal policy problem

The standard textbook approach to determining the optimal policy is to attach a Lagrangian multiplier, say γ_t, to the forward-looking constraint (3), to then find the first order conditions, and to finally determine the optimal behavior of the inflation and output by solving the resulting linear difference equation system under rational expectations (see Gali [2008, chapter 4], Walsh [2003] or Woodford [2003]).

In our analysis, we use recursive methods that also begin with Lagrangian multipliers as in the work of Marcet and Marimon [1998, 2011] on dynamic contracts and of Khan, King and Wolman [2003] on optimal monetary policy. Since the expectations of private agents are conditioned on the inflation realization, π_t, then the relevant *commitment multiplier* needs to be contingent on this realization, $\gamma_t(\pi)$, where the subscript t is short-hand for all history up through the start of the period.[13] Since the policy action must be taken before the inflation realization, we write it simply as a_t. Furthermore, a_t from now on refers to the policy action taken by the strong monetary authority (type $\tau = 1$) unless specified otherwise, since we are solving for the optimal policy assuming the current monetary authority is the strong type.

2.4 Two-stage optimal policy structure

The two stage nature of monetary authority and private sector decisions in our setting is reflected in a two stage nature of the recursive programming problem that we solve. Define the *interim value function* Ω via

$$\begin{aligned}\Omega(\rho_t, \eta_t, \varsigma_t, \pi_t) \quad = \quad & \min_{\gamma_t} \max_{x_t} \{u(\pi_t, x_t) \\ & +\gamma_t[\pi_t - \kappa x_t - \varsigma_t] - \eta_t(l_t \pi_t + (1-\rho_t)[\mu + \phi\varsigma_t]) \\ & +\beta E W(\rho_{t+1}, \eta_{t+1}, \varsigma_{t+1})|\eta_t, \varsigma_t, \pi_t\}\end{aligned} \tag{8}$$

with

$$\eta_{t+1} = \gamma_t. \tag{9}$$

representing the evolution of the *pseudo-state variable* η in terms of the commitment multiplier γ and with

$$\rho_{t+1} = b(\pi_t, \rho_t; a_1, a_2)$$

[13] See Appendix A and King, Lu and Pasten [2011]

being required by (6). Also define the *initial value function* W as

$$W(\rho_t, \eta_t, \varsigma_t) = \max_{a_t} \int \Omega(\rho_t, \eta_t, \varsigma_t, \pi_t) dF(\pi_t|a_t) \tag{10}$$

where $F(\pi|a)$ is the distribution of inflation conditional on a particular policy action.

We establish the appropriateness of this recursive system in KLP, using arguments briefly summarized in Appendix A, so that we focus here on its economic content. The policy action, a_t, must be made by the monetary authority without exact knowledge of its ultimate consequences for inflation so that the form of (10) is intuitive. That is, the optimal policy plan is one that maximizes the expected objective given the uncertain inflation outcome. After the realization of inflation, the monetary authority can take no direct action. However, the design of its optimal policy plan takes into account that there will be consequences of its future actions for how expected inflation responds to the actual inflation outcome. In turn, the expectations response governs how output responds to inflation given the forward-looking constraint (5). An outcome of the optimization in (8) is a pair of *contingency plans* for output

$$x(\rho_t, \eta_t, \varsigma_t, \pi_t) \tag{11}$$

and for the commitment multiplier

$$\gamma(\rho_t, \eta_t, \varsigma_t, \pi_t) \tag{12}$$

that is attached to (5). Note that the policy action does not enter directly in the decision problem (8) or the decision rules (11,12), but it does determine the distribution of inflation outcomes. Note also that the pseudo-state variable η_t could be replaced by $\gamma(\pi_{t-1})$, but we opt for the present notation as it allows a clear separation between the contingency plan $\gamma(\rho_t, \eta_t, \varsigma_t, \pi_t)$ and the manner in which the commitment multiplier serves as a state variable. Put concretely, given $\eta_t = \gamma_{t-1}$, other elements of history such as $\eta_{t-1}, \varsigma_{t-1}, \pi_{t-1}$ are irrelevant. The notation also is consistent with the general framework of Marcet and Marimon [1998, 2011]. The choice of the commitment multiplier γ is the way in which the recursive representation captures the management of expectations conditional on π_t.

2.5 Interaction of credibility and policy

We will see that the extent of policymaker credibility (ρ) will have major implications for the nature of optimal policy undertaken by a strong policymaker. At this stage, it is therefore useful to review the four model components where credibility enters. Doing so, we identify four channels of effect.

2.5.1 Effects of credibility on the trade-off

The inflation specification (3) implies that

$$\pi_t = \kappa x_t + \beta l_{t+1}[E_t a_{1,t+1}] + \beta(1 - \rho_{t+1})[\mu + \phi E_t \varsigma_{t+1}] + \varsigma_t$$

In terms of the trade-off between inflation and output that constrains optimal policy, there is a *level effect* on the trade-off $\beta(1 - \rho_{t+1})[\mu + \phi E_t \varsigma_{t+1}]$, and a *slope effect*, $\beta l_{t+1}[E_t a_{1,t+1}]$, with $l_{t+1} = \rho_{t+1} + \omega(1 - \rho_{t+1})$. Each of these influences the consequences of the current policy action a_{1t} or the future policy actions such as $a_{1,t+1}$. Generally, higher credibility reduces the level effect and raises the slope effect. With an automaton alternative policymaker ($\omega = 0$), credibility variable is evidently relevant for the slope ($l_{t+1} = \rho_{t+1}$). However, if there is a tag-along alternative policymaker ($\omega = 1$), then there is no slope effect because $l_{t+1} = 1$ always.

2.5.2 Evolution of endogenous credibility

The next two channels are reputation/learning effects which operate through

$$\rho_{t+1} = b(\pi_t, \rho_t; a_1, a_2) = \frac{\rho_t \psi(\pi_t; a_1)}{\rho_t \psi(\pi_t; a_1) + (1 - \rho_t)\psi(\pi_t; a_2)}$$

where $\psi(\pi; a)$ denotes the probability of observing π conditional on the policy action being a. A higher level of credibility ρ_t directly has a *level effect* on future credibility ρ_{t+1}.

The *marginal learning effect* of the action a_{1t} is more subtle, as it depends on the assumed relationship between a_2 and a_1. Our benchmark automaton assumption is that a_2 is invariant to a_1,. Under this assumption, a lower policy action serves to reduce the inflation outcome – at a given implementation error – and raise ρ_{t+1}. However, under our tag-along assumption ($\omega = 1$ implies $a_2 = a_1 + \mu + \phi\varsigma$), there is essentially no marginal learning effect as we explain further below.

2.5.3 Diagnostic and substantive model variations

In our research design, we initially introduced the tag-along alternative policymaker as a simple device for learning about the likely implications of strategic ("mimicking") behavior by an alternative policymaker: this specification gives the strong policymaker full leverage on expectations and removes the marginal learning effect of its policy action relative to the benchmark automaton alternative. However, in working through the results of various experiments below, we found that selectively varying the ω parameter in just one of the two equations above allowed us to diagnose the origins of policy response in expectations leverage and marginal learning separately.

2.6 Policy announcements and signalling games

So far, we have not been explicit about the policy announcement made by the monetary authority at the start of each period after viewing the shock ς_t and the realization of inflation from the prior period π_{t-1}. If the current monetary authority is the strong type, it will announce its planned action that was already chosen at time zero since the plan is ex-ante optimal and the strong type, by definition, has committed to that plan. If the current monetary authority is the alternative type, the current paper assumes that it will make the same policy announcement as the strong type. The rationale of imposing such an assumption is that the equilibrium outcome obtained under this assumption is consistent with the equilibrium outcome in an explicitly modelled signalling game in which both the strong type and the alternative type are strategic message senders and the private sector learns from the policy announcement about sender's type. A detailed study about the signalling equilibrium is beyond the scope of the current paper (since we do not model a strategic alternative type) but Lu [2012] in which such an equivalence result is established for a setup with a strategic alternative type.

3 Constant credibility case

We begin by exploring optimal policy when there is constant credibility and the alternative type is an automaton following a more inflationary policy rule detailed below. The results reported in this section set a benchmark for our endogenous credibility analysis in sections below and it also allows a comparison to the work of Schaumberg and Tambalotti [2007], who also study optimal policy in a setting where agents are skeptical about the degree of policymaker commitment. In their analysis, the monetary authority recognizes that it will be replaced with a probability that is also known by private agents. By contrast, in our context, the monetary authority knows that it will be present forever, but also recognizes that private agents are skeptical about its identity and behavior.

3.1 Explicit solution

While the recursive approach is general enough to be applied to economies without a quadratic momentary objective (2) or a linear forward-looking constraint (3), these additional assumptions allow us to derive an exact quadratic solution for the value functions Ω, W and an exact linear solution for the decision rules a, γ, x.

As shown in Appendix B, the solutions can be derived using the first order conditions to the optimization problems (8) and (10). We write the solutions in terms of slope coefficients θ and stationary values $\underline{a}, \underline{\gamma}, \underline{x}$: as is standard in rational expectations models, these solution coefficients depend nonlinearly on the structural coefficients, which here includes the constant level of credibility ρ. The optimal policy plan is

$$a_t = \underline{a} + \theta_{a\eta}(\eta_t - \underline{\eta}) + \theta_{a\varsigma}\varsigma_t \tag{13}$$

and the contingency plans are

$$\begin{aligned} \gamma_t &= \underline{\gamma} + \theta_{\gamma\eta}(\eta_t - \underline{\eta}) + \theta_{\gamma\varsigma}\varsigma_t + \theta_{\gamma\varepsilon}\varepsilon_t & (14) \\ x_t &= \underline{x} + \theta_{x\eta}(\eta_t - \underline{\eta}) + \theta_{x\varsigma}\varsigma_t + \theta_{x\varepsilon}\varepsilon_t & (15) \end{aligned}$$

with (9) implying that $\underline{\gamma} = \underline{\eta}$ and

$$\eta_{t+1} - \underline{\eta} = \gamma_t - \underline{\gamma} \tag{16}$$

As noted above, the decision rule coefficients are all functions of ρ, the extent of credibility of the plan. Thus, with constant credibility and implementation error, we have a generalization of the closed form decision rules familiar from review articles and textbooks. In fact, the optimal policy plan has a certainty equivalence property: the introduction of implementation error does not affect the coefficients governing optimal transition dynamics ($\theta_{a\eta}$, $\theta_{\gamma\eta}$) or shock responses ($\theta_{a\eta}$, $\theta_{\gamma\varsigma}$).

Table **3.1** shows the formulas governing these coefficients and Table **3.2** shows numerical solutions for these coefficients at three different values of credibility ($\rho = 0, .5, 1$). The *fifty-fifty* case in which $\rho = .5$ and $1 - \rho = .5$ is a particularly useful intermediate case, as it is midway between the no credibility and full credibility assumptions in the literature. It will play a key role in our discussion of exogenous credibility in this section and of endogenous credibility in later sections.

3.2 Implications for a policy rule

The policy plan (13) contains the unobservable multiplier η, so that it is useful to represent the policy action solely in terms of observables (as done in the survey by Clarida, Gali and Gertler [1999] and advocated as a general strategy by Gianonni and Woodford [2004]). This representation, which we'll call the simple rule, is

$$\begin{aligned} a_t &= \underline{a} + \theta_{\gamma\eta}(a_{t-1} - \underline{a}) + \theta_{\gamma\eta}\varepsilon_{t-1} \\ &\quad + \theta_{a\varsigma}\varsigma_t + [\theta_{a\eta}\theta_{\gamma\varsigma} - \theta_{\gamma\eta}\theta_{a\varsigma}]\varsigma_{t-1} \end{aligned} \tag{17}$$

using the same coefficients discussed above. As a reference point for analysis below, Figure **3.1** graphs the values of these coefficients against ρ for a reference set of structural parameters, for a quarterly time interval. Appendix C discusses this calibration in detail.

3.3 Inflation and stabilization bias

The coefficients in the simple rule (17) reflect the core results of the existing literature on optimal policy in the New Keynesian model at the endpoints $\rho = 1$ and $\rho = 0$.

First, there is an inflation bias from lack of commitment as in many macroeconomic models, reflected in the coefficient $\underline{a} > 0$ when $\rho = 0$. In this situation, the policy rule only reacts to contemporaneous variables. Second, there is a stabilization bias in that the coefficients $\theta_{a\varsigma}$ and $\theta_{x\varsigma}$ are larger when $\rho = 0$ than when $\rho = 1$, reflecting the fact that the monetary authority with full credibility can better stabilize inflation and output in response to the temporary inflation shock ς.[14] This extra stabilization stems from the ability of the monetary authority with full credibility to reduce the expected inflation, which provides an additional channel to offset the effect of a positive inflation shock compared to the monetary authority with zero credibility. Notice that the policy responds to the lagged inflation shock negatively when $\rho = 1$, whereas the response coefficient is zero when $\rho = 0$.

For our reference parameter values, the inflation bias is $\mu = 1\%$ or about four percent per year. Under commitment, the optimal policy response to ς shock is $\theta_{a\varsigma} = 1.04$. By contrast, without commitment, the response is $\theta_{a\varsigma} = 1.98$. The output consequences of these shocks are $\theta_{x\varsigma}(\rho = 1) = -10.37$ and $\theta_{x\varsigma}(\rho = 0) = -19.82$. Further, there is also stabilization bias associated with the response to implementation errors, as discussed further below.

With ρ changing continuously between 0 and 1, the extent of these biases changes smoothly. At our fifty-fifty reference case of $\rho = .5$, note that steady state inflation is positive and, in fact, $\underline{a}$ is higher than $.5 * \mu = 0.5\%$. Thus, our partial commitment model works differently from that of Schaumberg and Tambalotti [2007], where both the monetary authority and private agents correctly understand that there is a fixed exogenous probability of policymaker replacement each period and notably an optimal policy chooses a zero inflation rate if policymaker replacement does not occur (which would imply $\underline{a} = 0$ in our framework).

In this section, our policymaker's calculus is based on his knowledge that he will always be in place, but also that there will always be private sector skepticism about whether inflation will evolve according to his plan. He therefore more than accommodates the adverse inflation shift in inflation expectations, $(1-\rho)\mu$. The top panel in Figure **3.1** shows that this accommodation is a general result for all levels of credibility. To understand why, notice that the authority has an expected inflation $E_t \pi_{t+1} = \rho E a_{t+1}|(s_t) + (1-\rho)\mu$ and a Phillips curve trade-off

$$\pi_t = \beta \rho E a_{t+1}|(s_t) + \beta(1-\rho)\mu + \kappa x_t + \varsigma_t$$

since the automaton assumption about the alternative type implies that $l = \rho$. It faces a higher intercept of that trade-off ($\beta(1-\rho)\mu$ rather than 0 without skepticism) and also a worsened slope in terms of the effects of expectation management ($\beta\rho$

[14] We are using stabilization bias here in the sense of Svennson (1997), Clarida, Gali and Gertler (1999) and Woodford (2003).

rather than β without skepticism).

With these two alterations in its constraints, the monetary authority's policy action in a steady state can be shown to be

$$\underline{a} = \{\frac{(1+q(1-\beta\rho))}{1+q(1-\rho)(1-\beta\rho)}\}[(1-\rho)\mu]$$

with $q = h\kappa^{-2} > 0$, drawing on results in Appendix B.[15] The term in brackets is the level shift in expectations. The term in braces imbeds the effects of the slope $\beta\rho$ in the Phillips curve: the effect of expectation management is zero when $\rho = 0$ and is increasing in ρ. Hence, there is accommodation of expected inflation and the accommodation is more important when there is greater skepticism, even though the policymaker has full commitment.

3.4 Comparative dynamics

In terms of the dynamic behavior, we stress three points in the context of the simplified form of the rule (17). We start with two properties that the solution shares with the full commitment case. We explore these properties within three experiments portrayed in Figures **3.2** through **3.4** that capture optimal policy by a new monetary authority starting up (transitional dynamics), the response to an implementation error (ε), and the response to a price shock (ς).

3.4.1 Persistence in policy actions

Under full commitment, the persistence in monetary policy actions comes about to manage the evolution of actual inflation and current expected inflation in a desirable

[15] Appendix B shows

$$\underline{a} = \frac{(1-\rho)}{1+q(1-\rho)(1-\beta\rho)}[q\kappa x^* + q\beta(1-\rho)\mu]$$

with $q = h\kappa^{-2} > 0$. This expression captures three ideas. First, there is no steady-state inflation with perfect credibility ($\rho = 1$). Second, there is inflation bias when there is no credibility. To make this internally consistent, we require that

$$\mu = \underline{a}(\rho = 0) = \frac{q\kappa x^*}{1+(1-\beta)q} > 0$$

Further, with this condition imposed, the result in the text is direct.

manner, recognizing that commitment requires that the monetary authority respects past inflation expectations. It is therefore natural that the coefficient attached to a_{t-1} is precisely that which governs the optimal dynamics of the commitment multiplier, $\theta_{\gamma\eta}$. The second panel in Figure **3.1** shows that $\theta_{\gamma\eta}$ increases with ρ , i.e., with the extent of credibility.

Figure **3.2** displays the inflation policy that would be followed by a new monetary authority without the pre-existing commitments. As well known, the New Keynesian model implies that there should be an initial interval of high, but declining, inflation: this anticipated reduction in inflation stimulates real economic activity, which is desirable because steady-state output is inefficiently low ($x^* > 0$). As it is also well known, zero long-run inflation is optimal in the NK model.

With partial credibility, as shown in Figure **3.2**, there is also an interval of high but declining inflation as the economy converges to a positive long-run inflation rate. However, this interval is smaller in scale and shorter in duration due to the smaller degree of leverage that the authority has over expected inflation. Put alternatively, a smaller value of the $\theta_{\gamma\eta}$ coefficient with $\rho < 1$ means that the inflation action is less serially correlated when the level of credibility is lower.

3.4.2 Optimal response to implementation errors

Under full credibility, the response to implementation errors is governed by the same coefficient, $\theta_{\gamma\eta}$, and for a subtle version of the same reasoning. Last period, with inflation determined stochastically by $\pi_{t-1} = a_{t-1} + \varepsilon_{t-1}$, the only control that the committed monetary authority had over the response of output to the implementation shock was via inflation expectations, $E_{t-1}\pi_t = Ea_t|(\pi_{t-1}, s_t)$. Hence, it is natural that the monetary authority's current response to past monetary actions and past monetary policy errors reflects its desire to manage the response of expected inflation to actual inflation. Figure **3.3** plots the impulse responses for various values of ρ to a one-percent implementation error at date 1. With full credibility, the results look just like the transition dynamics just discussed: unexpectedly high inflation arising from an implementation error is optimally followed by an interval of higher-than-average inflation.

It turns out that the equivalence between startup and implementation error responses carries over to situations with alternative constant levels of credibility, although the strength is diminishing in ρ. As in the discussion of steady-state inflation, the weakened response – less persistent transitional dynamics and less persistent policy response to implementation shocks – reflects the fact that the policymaker sees only part of expected inflation responding to his policy actions. Continuing the discussion of the response to an implementation error last period, the policymaker faces inflation expectations $E_{t-1}\pi_t = \rho E a_t|(\pi_{t-1}, s_t) + (1-\rho)\mu$ and cannot as effectively manage these to offset ε_{t-1} when $\rho < 1$. When $\rho = .5$, the parameter $\theta_{\gamma\eta} = .34$, more than half of the way to its $\rho = 1$ value of .54. This extra response of current policy action to an implementation error last period reflects the decreased persistence in future monetary policy's responses due to imperfect credibility.

3.4.3 Serially correlated inflation shocks

We finally consider how partial credibility affects the response to inflation shocks, current and lagged. When there is no commitment, Clarida, Gali, and Gertler [1999] show that there is an inflation policy which depends only on ς_t since the policymaker has no control over expectations. Hence, in Figure **3.1**, the coefficients on a_{t-1}, ς_{t-1} and ε_{t-1} are all zero when $\rho = 0$. The coefficient on ς_t is about 1.98. When there is full commitment, optimal policy is a form of "flexible price level targeting"[16]. Accordingly, the coefficients in Figure **3.1** on the current and lagged inflation shock terms must be equal in absolute value and of opposite sign when $\rho = 1$, so that there is no long-run effects of these shocks on the price level under full commitment.

When $\rho = .5$, as with other intermediate values of ρ, there is a flexible inflation target rather than a price level target, as can be seen by the fact that the coefficients on ς_t and ς_{t-1} are no longer equal in magnitude and opposite in sign. In our fifty-fifty case, the coefficient on ς_t is 1.83 (as opposed to a full commitment value of 1.04) and the coefficient on ς_{t-1} is $-.91$ (as opposed to a full commitment value of -1.04). Given the response coefficient on the lagged action is .34, there is thus a protracted response of inflation to serially correlated inflation shock.

Figure **3.4** plots the impulse responses to a 1% inflation shock at date 1 with persistence .9. This inflation shock has a contractionary effect at all credibility levels. The path of inflation action when $\rho = 0$ reflects the persistence of the shock. The response of inflation action when $\rho = 1$ is first positive and then negative, reflecting the optimality of "flexible price level targeting".

[16] See King and Wolman [**1999**] and Woodford [**2003**] for early statements of the case for price stability made by this class of models and Gali [**2008, chapter**] for a textbook presentation.

3.5 A tag-along alternative policymaker

We now consider a tag-along alternative policymaker. Under this assumption, the strong policymaker gains complete leverage over expected future inflation with l increasing from ρ to 1. However, there policymaker is still confronted by the level effect on expected future inflation, which takes the form $(1-\rho)(\mu+\phi\varsigma_t)$. As a consequence, the trade-off between inflation and output facing the policymaker takes the following form

$$\begin{aligned}
\pi_t &= a_t + \varepsilon_t \\
&= \beta E a_{t+1}|(s_t) + \beta(1-\rho)[\mu + \phi\varsigma_t)] + \kappa x_t + \varsigma_t \\
&= \beta E a_{t+1}|(s_t) + \{\beta(1-\rho)\mu + \kappa x_t\} + \{1 + \beta(1-\rho)\phi\}\varsigma_t
\end{aligned}$$

Relative to the full credibility case, there are thus two consequences that can be interpreted using the reasoning standard in the literature since certainty equivalence is assured in this section. First, the policy maker acts as if he faces a different value of x^*. However, the optimality of zero long-run inflation does not depend on the value for x^*, so that it is optimal when there is only a level effect on expected inflation. Optimal long-run output is

$$\underline{x} = -\frac{\beta(1-\rho)\mu}{\kappa}$$

Second, the policymaker acts as if he faces a more volatile structural inflation shock, $\{1+\beta(1-\rho)\phi\}\varsigma_t$. Thus, the optimal pattern of dynamic response is simply a re-scaled version of the full credibility case in Figure **3.4**.

Thus, there is no inflation bias when the alternative policymaker follows the tag-along rule, irrespective of the extent of credibility. However, there is a persistent effect on output and there are implications for stabilization bias.

3.6 Summarizing the level and leverage effects

In the fixed credibility case, we have identified two effects of lower credibility on the output-inflation trade-off, both of which operate through expected inflation: one is a level effect that raises the rate of inflation at a given output gap and the other is a slope effect that reduces the impact of expected future policy actions on current inflation at a given output gap. The manner in which these two mechanisms affect the nature of optimal monetary policy depend on how private agents perceive the alternative monetary authority as behaving. If it is an automaton that is more inflationary, then the optimal policy is accomodative after an initial startup period: the steady state inflation bias is higher with lower credibility, while the extent of the startup component is smaller and less protracted. If it is a tag-along type that is more inflationary, then the strong authority does not suffer a loss in leverage with lower credibility: the long-run inflation rate is zero and the startup component identical for all levels of credibility (except zero). The real consequences of imperfect credibility depend as well on the nature of the alternative: these are largest when there is a tag-along alternative and when credibility is low.

4 Evolving reputation: a benchmark case

We now turn to exploring how a strong monetary authority – confronted with private sector skepticism about its commitment ability but recognizing that there is learning – will design its optimal policy. As in the prior section, the strong authority understands that the private sector believes that policy may alternatively be undertaken by an automaton that behaves according to $a_2 = \mu + \phi \varsigma_t$,where the μ and ϕ parameters are set to match the equilibrium behavior that would arise if there were no commitment in monetary policy. Colloquially, the private sector is unsure about whether it is facing a discretionary central bank or a committed one. As above, the extent of the inflation bias is 4% per year (so that $\mu = .01$).

4.1 General points about learning

Before getting into the details, it is important to emphasize two general points about our simple and very standard Bayesian learning specification. First, realized inflation (π) is equal to the policy action (a) plus an implementation error (ε),

$$\pi_t = a_t + \varepsilon_t$$

with ε_t being distributed $N(0, \sigma^2)$. We assume that there is a set of actions A and a set of outcomes Π such that the truncated normal approximation is reasonable at all points in the set Π. Accordingly, Bayesian learning takes the form

$$\rho_{t+1} = b(\pi_t, \rho_t; a_1, a_2) = \frac{\rho_t \psi(\pi_t; a_1)}{\rho_t \psi(\pi_t; a_1) + (1-\rho_t)\psi(\pi_t; a_2)}$$

where $\psi(\pi; a) = \frac{1}{\sigma}\exp(-\frac{(\pi-a)^2}{2\sigma^2})$.[17] Equivalently, the learning specification can be written as

$$\begin{aligned} b(\pi_t, \rho_t; a_1, a_2) &= \frac{\rho_t}{\rho_t + (1-\rho_t)\exp(\frac{[(\pi_t - a_{1t})^2 - (\pi_t - a_{2t})^2]}{2\sigma^2})} \\ &= \frac{\rho_t}{\rho_t + (1-\rho_t)\exp(\frac{[-2\varepsilon_t(a_{1t} - a_{2t}) + (a_{1t} - a_{2t})^2]}{-2\sigma^2})} \end{aligned}$$

In the absence of realized implementation errors ($\varepsilon = 0$), there will be a growth in ρ that takes place more rapidly ($|a_{1t} - a_{2t}|$ is larger) or the degree of randomness is smaller (σ).

In section 2, we wrote the learning mechanism as $\rho_{t+1} = b(\pi_t, \rho_t; a_1, a_2)$ because we want to highlight the idea that reputation/credibility is like a capital good. From this perspective, as illustrated in the top panel of Figure **4.1**, a lower inflation action induces more rapid learning so that low inflation is like investment. The top panel shows that learning can be very fast, if there is a wide gap between the policy actions even if there is quite a lot of implementation error. If the extent of inflation bias is 1% quarterly, as in calibration which we use, and if the standard deviation is also 1% then a zero inflation policy produces a very fast pace of learning: ρ climbs from .2 to about .9 in just 8 quarters (note that the paths drawn are ones that assume a zero implementation error realization, the mode of the implementation error).

Further, empirical models with learning like those of Baxter [1985], Schorfheide [2005] and Milani [2007] stress that a positive implementation error $\varepsilon > 0$ will lower the likelihood of the "good regime" ($a_1 < a_2$) as well as that this effect is stronger if a shock occurs when ρ is very small than when ρ is very large. The second panel of Figure **4.1** illustrates this effect. However, in contrast to such empirical analyses, our model assumes that the action level a_1 can itself adjust to the state of the economy (including ρ).

[17] We drop the factor $(2\pi)^{-1}$ in the normal pdf from the front of ϕ to avoid confusion with the inflation rate.

4.2 Comparative dynamics

We now consider three responses to initial conditions and shocks, focusing principally on the policy action but also providing some information about the nature of macroeconomic outcomes.

We explore the consequences of a policymaker having an inherited reputation ρ_0 at five alternative values: $0, .25, .5, .75, 1$. These are the three nodes for our polynomial approximation and two equally spaced intermediate values. In each of the four panels of Figures **4.2-5**,there are 5 different paths, corresponding to the five different initial initial conditions on ρ.

We consider start-up dynamics, response to an implementation error, and response to a serially correlated inflation shock in turn as we did in section 3 for the exogenous credibility case. In each of the figures in this section, the full credibility ($\rho = 1$) and no credibility ($\rho = 0$) solutions are exactly the same as in the corresponding figure in the previous section. For example, the start-up dynamics under full credibility and no credibility in Figure **4.2** are exactly the same as in Figure **3.2**: this coincidence arises since reputation is not an endogenous variable from these initial conditions. The full credibility path is indicated by ('*') and the no credibility path by ('o').

The dynamics arising with endogenous credibility are remarkably different from those just considered with exogenous, constant credibility in the prior section. After we document these differences using Figures **4.2-4.4**, we turn to decompositions that are revealing about the key sources of these differences in section 5.

4.2.1 Start-up dynamics

One celebrated implication of the New Keynesian model under optimal policy is that there is an initial period of high but declining inflation, during which the economy experiences high real economic activity. We saw in section 3 that this implication was sustained with exogenous and constant credibility. We now investigate the robustness of this implication to the assumption that credibility is incomplete, but endogenous.[18] Panel A of Figure **4.2** shows the sequence of monetary policy actions, a, taken by the committed policymaker at each date, under the assumption that no implementation errors actually arise ($\varepsilon = 0$) and that no structural inflation shocks occur ($\varsigma = 0$), which is the same assumption which was made in Figure **3.2** when

[18] Symbol references are 1('*'),.75('$\triangle$'),.5('$\diamond$'),.25('$\bigtriangledown$'),0('o').

reputation was held constant. The subsequent panels display expected inflation e (panel B), reputation/credibility ρ (panel C) and real output x (panel D).

Fifty-fifty: Consider first a monetary authority which starts with $\rho = .5$, i.e., private agents believe that there is a 50% likelihood that they are facing a policymaker of either type. With exogenous credibility in Figure **3.2**, we saw that the optimal policy involved an initial inflation rate of about 3.6% which was gradually reduced to a long-run rate of 2.9%.

The results with endogenous credibility are sharply different, both in the short-run and the long-run, as can be see by looking at the '$\diamond$' path in panel A of Figure **4.2**. The monetary authority chooses to eliminate inflation immediately (its initial action a is close to zero) and to follow up with deflationary actions: these are taken so as to build its reputation, which rises sharply in Panel C (reaching $\rho = .9$ within a year). Its ability to invest in reputation means that it asymptotically chooses zero inflation, in contrast to its choice of a positive inflation rate in the exogenous reputation case that was displayed in Figure **3.2**.

Thus, in the fifty-fifty case, endogenous credibility overturns both key implications of the NK model with optimal policy choice in the presence of skepticism: there is no start-up inflation and there is no long-run inflation.

Turning to the details of the transitional dynamics, we see that expected inflation is dramatically affected by the endogeneity of reputation, since private agents understand that a committed authority will take tough actions: panel B shows that it is about 1% in the first period, zero in the second period, and turns negative thereafter. But, with expected inflation always above actual inflation and with the extent of this difference evolving over time, there is a recession that is initially quite deep as shown in panel D (the output is about -6%, with a gradual recovery taking place over a year). The persistently low level of output reflects the tough actions taken by the monetary authority and skepticism that private agents hold toward these actions, which are resolved only after a year or so.

Stronger initial reputation: Similar outcomes arise when the monetary authority has a stronger initial reputation ($\rho = .75$, illustrated with a '$\triangle$' in Figure **4.2** and others below). With less of an investment in reputation to be accomplished, the optimal policy actions are somewhat less restrictive than those the fifty-fifty case, but muted in magnitude. As in the fifty-fifty case, the desirability of investing in reputation leads to the elimination of the initial interval of "start-up" inflation that arises under optimal policy with credibility.

Weaker initial reputation: A very different dynamic path arises when the monetary authority has a weaker initial reputation ($\rho = .25$; paths indicated by ('$\bigtriangledown$'): there is an initial interval of inflation, resembling the start-up solution for several periods. But the motivation for the gradual reduction in inflation appears to us to be very different in these two situations. For the policy maker with full commitment, a modest boom is created by having the expected rate of inflation below the current rate of inflation (see the '*' path in Figure **4.2**).

By contrast, for the policymaker with a weak initial reputation ($\rho = .25$), inflation expectations are relatively insensitive to his policy action. Hence, an aggressive reduction in inflation – say, a "cold turkey" zero inflation solution – is simply too costly in terms of output. Even under the optimal policy, a deep recession occurs, with an 8 percent output gap arising for the first year.

In the standard analysis of policy without credibility, the monetary authority is unwilling to fight inflation – when it is at its equilibrium value – because there would be output losses from doing so. It takes inflation expectations as unaffected by its policy actions. The monetary authority with weak initial reputation does not suffer from a commitment problem in our framework. However, it does face skepticism concerning its inflation plans manifested in a reduced leverage on expectations, i.e., it has imperfect credibility. Accordingly, its optimal commitment policy is to gradually reduce inflation as it balances output losses and credibility building.

4.2.2 Implementation error

We now turn to the effect of a one-time implementation error, ε_t, at date $t = 0$. There are several points to keep in mind as we examine the responses in Figure **4.3**.[19]

First, the implementation error occurs after the policy action a_t so that there is no initial period policy response in panel A or in the predetermined reputation state variable ρ_t in panel C. Second, the effect of a positive implementation error in the full credibility case of Section 3 (Figure **3.3**) is to increase output if expected inflation is held fixed: this benchmark again is displayed as the $\rho = 1$ path in Figure **4.3** (the '*' path). With full credibility, the optimizing monetary authority chooses to have expected inflation increase, thus partially mitigating the impact effect on

[19] Note that all of the implementation error results are to be interpreted as one would an impulse response, in that they represent deviations from the paths shown in Figure **4.3**.

output and, in effect, smoothing the shock's effect by raising output and inflation in subsequent periods . That is, in the full credibility case, the positive implementation error induces a higher path of inflation, resembling the start-up dynamics. Third, by contrast, with zero credibility, the monetary policy authority has no influence on expectations and its future behavior does not respond to this one time inflation shock. Hence, the date t=1 output effect is maximized (at $\varepsilon_t/\kappa = 25\%$ annually and 6.25% quarterly) and there is no persistence in inflation or real activity.

Turning to the imperfect credibility results in Figure **4.3**, we again use the full credibility case ('*') and the no credibility case ('o') as reference points.

Fifty-fifty: Focusing on the monetary authority with initial reputation of $\rho_0 = .5$ (the '$\diamond$' paths in the figure), we see that optimal policy response allows expected inflation to rise in period 0, in part by increasing a_1, so that the output effect of the implementation error is muted on impact just as in the full credibility case. However, there are new elements at work with $\rho_0 = .5$: the implementation error shock causes a decline in ρ_1 so that rebuilding of reputation will be necessary in the future. Because credibility has deteriorated, the policymaker experiences some of the same challenges in terms of output that we described for the transitional dynamics from the fifty-fifty initial condition: there is a small accommodation to expected inflation (change in a that is positive) at date 1, after which the rebuilding of reputation commences. During this investment period, there is a negative inflation response, so as to allow the monetary authority to better distinguish himself from the alternative type.

Stronger initial reputation: The authority with initial credibility of $\rho = .75$ also allows expectations to rise in response the implementation error (the level of a is positive in panel A of Figure **4.3**). As a consequence of the show, he faces a deterioration of reputation as shown in panel D of Figure **4.3**. But, since the decline is smaller and because his action at date 1 is slightly tougher, the rebuilding process for reputation begins immediately in date 1.

Weaker initial reputation: The policymaker with initial credibility of $\rho = .25$ behaves very differently. As shown in panel A of Figure **4.3**, he allows the implementation error in period 0 to generate a positive movement in his action for the first four quarters after the shock. In terms of magnitudes, the one percent inflation shock at date 0 causes the policy action to be elevated by an average of one-half percent over the first four quarters. As in our analysis of Figure **4.2**, the policymaker with weaker initial reputation is accommodative because he faces a very major upward shift in expected inflation, due to the decline in reputation. Tougher policy actions would have adverse output effects that he seeks to avoid. Note that he is already, in Figure **4.2**, undertaking a gradual disinflation with substantial output costs. Consequently, the monetary authority invests less in reputation building for about a year when faced with a positive implementation error, prolonging the relatively lengthy reputation-building process in panel C of Figure 4.2.

4.2.3 Serially correlated inflation shocks

To be added.

5 Leverage, Learning, and Strategic Behavior

Endogenous credibility can have substantial implications for the behavior of inflation and real activity under optimal policy, given the results reported in the previous section. To understand why, we now explore alterations in structural elements of our endogenous credibility model including analysis of "tag-along" behavior of the alternative policymaker.

As discussed at the end of section 2, credibility interacts with policy within two components of our model. First, credibility affects the influence of expected future policy in the inflation equation (7). Relative to the benchmark studied in the last section, we can restore the complete leverage that the monetary authority has by adjusting the value of ω to one in this equation. For concreteness, call the value of ω in this equation ω_p. Second, the current policy action affects the evolution of credibility in the Bayesian learning rule, but this effect is shut-down if ω is set to one. For concreteness, call the value of ω in this equation ω_b. Finally, if we set $\omega_p = \omega_p = \omega = 1$, then we obtain full tag-along behavior. Using this approach, we thus study the three cases displayed in Table **5.1**. To keep the analysis manageable, we restrict attention to the start-up dynamics.

5.1 Suppressing the effect of policy on learning

Given the analysis of sections 4 and 5, it is natural to begin by examining a variant of our basic endogenous credibility model that rules out the effect of policy actions on learning, which can be accomplished by setting the parameter $\omega_b = 1$ while keeping the parameter $\omega_p = 0$.[20] This is closely related conceptually to the exogenous credibility analysis of section 3 but there is one crucial difference: while credibility is unaffected by the policy action, it is not constant over time but rather evolves according to the difference equation that is Bayes' law. Panel C of Figure **5.1** shows that the evolution of reputation depends substantially on the initial condition: it is quite fast for $\rho_0 = .75$ (so that ρ is nearly one after two years) and quite slow for $\rho_0 = .25$ (so that ρ attains .4 only after four years).

There were two key aspects of our section 4 analysis of optimal policy with an automaton alternative and endogenous credibility, relative to the constant exogenous credibility models of section 3: an elimination of the "start up" interval of high inflation and the asymptotic elimination of inflation. This diagnostic experiment shows that the first of these does not occur when the effect of policy actions on learning is eliminated. Policy is always more inflationary than the full credibility solution, being most inflationary for low credibility (this finding harks back to the constant exogenous credibility case of section 3).

But, so long as $\rho_0 > 0$, reputation will asymptotically approach 1. Hence, the zero long-run inflation implication is obtained in all cases. In the fifty-fifty case, the optimal policy is to reduce inflation from about 2.8% to about 0% over roughly a year with a inflation falling by roughly the same amount each quarter. Relative to the optimal policy path displayed in Figure **4.2**, the elimination of learning means that (i) there is a slower reduction in inflation; and (ii) there is a never a deflation. Put alternatively, the diagnostic experiment in this subsection confirms our earlier assertion that it is policy concern about learning which makes it aggressive in Figure **4.2**, both in terms of the speed of inflation elimination and the desirability of deflation as part of the optimal policy.

5.2 Eliminating the leverage loss from imperfect credibility

We next consider the reverse diagnostic experiment: eliminating the leverage loss from imperfect credibility (setting $\omega_p = 1$ so that $l_t = 1$ in every period) but maintaining the learning effect of section 4 (setting $\omega_b = 0$).[21]

In isolation, strengthening the monetary authority's leverage on expected inflation

[20] This is experiment 81 housed in \KLP April 2012\Codes\EndoCredibility\Experiment_81_IRtd
[21] This is experiment 82 housed in \KLP April 2012\Codes\EndoCredibility\Experiment_82_IRtd

makes it more desirable for the policy authority to have a gradual reduction in inflation. (Recall from section 3 that a permanent increase in credibility led to a greater initial inflation rate – relative to the relevant steady state – and a more measured reduction in inflation). Looking at the $\rho_0 = .25$ optimal policy path in Figure **5.2** and then comparing it to that in Figure **4.2**, notably, we see that greater expectations leverage leads to greater inflation in the early stages of the plan – about 2% rather than just over 1% – as well as a more rapid movement to an interval of deflation and it is somewhat more severe with increased leverage. The greater leverage in Figure **5.2** does lead to smaller output losses during the transition to price stability, but it does not eliminate these disinflation costs because there remains the level effect of imperfect credibility on the trade-off between inflation and output.

Taking the two of our diagnostic experiments together, we conclude that the learning mechanism is the main structural feature which leads to the nature of optimal policy within the benchmark endogenous credibility model of section 4, which features an automaton alternative.

5.3 Tagalong behavior

Much of the prior literature on reputational mechanisms in macroeconomic policy focuses on the incentives that these provide for a weak policymaker to behave in the same manner as a strong one. For example, building on the "chain store" results of Kreps and Wilson [1982] and Milgrom and Roberts [1982], Backus and Driffill [1985a,b] and Barro [1986] showed that a weak policymaker can be led to adopt the zero inflation policy of an automaton strong type until the later stages of a finite horizon game.[22]

We close our analysis of optimal policy design with imperfect credibility by supposing that the alternative monetary authority follows a tag-along behavioral rule, of the form $a_2 = a_1 + \mu + \phi\varsigma_t$ so that the inflation and stabilization biases arising from lack of commitment are adjusted one-for-one with the policy action of the strong authority. This ad hoc rule is chosen to represent, in a simple and transparent manner, the potential implications of policy mimicking as described in the earlier literature.

[22] A recent development in reputation literature by Cripps, Mailath and Samuelson (2004) shows that introducing imperfect monitoring in an infinite-horizon game undermines the incentive of the weak type to mimic in the long run. However, mimicking can still be a short-run phenomenon. In addition, the long-run mimicking behavior can be restored if the type of the long-lived player is governed by a stochastic process. See for example, Mailath and Samuelson (2001).

A more complete analysis of this topic will require developments along lines that we are pursuing in companion work.

As with our exploration of the automaton alternative in section 4, we consider three cases: the startup dynamics plus the responses to implementation errors and structural inflation shocks.[23]

5.3.1 Startup dynamics

Figure **5.3** shows that the results of policy mimicking by the alternative can be dramatic for the strong monetary authority: at all levels of credibility, there is a gradualist policy for inflation. The early stages of the disinflation path contain the two elements described by Ball [1994, 1995] and stressed by Goodfriend and King [2005] in their analysis of the Volcker disinflation: in real terms, there is a stimulative credible disinflation effect and a contractionary incredible disinflation effect. For our lower level of initial credibility ($\rho_0 = .25$), the two effects cancel out in the initial period, with the economy subsequently displaying declining inflation and an intensifying recession.

5.3.2 Implementation error

Figure **5.4** shows that policy mimicking of the tag-along variety leads to a protracted increase in the path of the inflation action when there is a positive policy implementation error. At all levels of credibility, there is an initial stimulative effect. But for the low credibility initial condition, the initial positive output response is followed by a negative one.

to be completed

5.3.3 Inflation shock

to be added

5.4 A closer comparision of the specifications

In our benchmark specification (section 4) and in the three specifications explored in this section (no learning, full leverage, and tag-along), we have seen a variety of

[23]NOTE: these are the results of Experiment 78, with the figures produced using `GroupPlot_endo_cred_tagalong (May 8 2012).tex`

behavior within specification as we varied the initial condition or perturbation. It is useful to take a closer look at how the specifications compare, with an eye to further determining the relevant channels of effect. To that end, Figure **5.6** shows the transitional dynamics from $\rho_0 = .25$ and Figure 5.6 shows the transitional dynamics from $\rho_0 = .5$ (the results from $\rho_0 = .75$ are quite similar to those from $\rho_0 = 0.5$).

Fifty-fifty: The benchmark path is the solid line in Figure 5.6. As discussed above, this involves an immediate drop to zero inflation followed by an interval of negative inflation. It also involves a loss of output that is initially quite substantial (-5%).

Moving from this benchmark to the specification with full leverage (similar to Figure 5.2 and with a path marked with square plotting symbol), we see that there is even tougher policy: negative inflation begins right away, but because there is greater control over expectations, there is a smaller and less persistent output loss.

The two specifications that suppress learning effects – the no learning and the tag-along cases – involve gradual disinflation. There is a smaller output loss for each than under the benchmark case, reflecting this more moderate policy path. But both also feature a triangular recessionary path for real activity as a result of the intensifying anti-inflation policy and the sluggish path of expectations.

Weaker initial reputation: More dramatic differences emerge when the initial condition is $\rho_0 = .25$. The benchmark includes an initial inflation rate of 1% and an initial output gap of -8%, with a subsequent period of about a year during which inflation falls to less than -1% while output stays depressed. With an initial condition of $\rho_0 = .25$, there is a initial level effect on expected inflation of (1-.25)*4%=3% that declines only as credibility rises and the monetary authority's leverage related to ancipated future policy actions is relatively weak. Hence, the increase in leverage afforded by one of our diagnostic experiments is quite important. With full leverage, the authority chooses an inflation action path that starts higher and declines more sharply, cushioning the path of output substantially.

Turning to the two specifications which suppress learning, there are gradualist outcomes as there were with higher initial credibility. However, as the tag-along case features full leverage of policy actions on expectations, while the direct no learning case involves weak leverage, there are now visually apparent distinctions between the tag-along (square) and no learning (hexagram) paths. With full leverage, the tag-along authority chooses an inflation action path that starts higher and declines more sharply than the policy path chosen by its no-learning counterpart. One notable feature is that output is reduced to a greater extent for the the tag-along authority than for its no-learning counterpart. Viewed from the standpoint of the section 3 fixed credibility case, this is not too surprising: there, an authority with full leverage chose to bring about zero inflation for every credibility level, even though there were permanent output reductions in a size that was increasing in output. Further, the tag-along policymaker faces a particular challenge. With full leverage, it would like to have an inflation path that starts positive and declines, so as to induce a stimulation. However, if it does so, inflation expectations are higher at any credibility level because private agents correctly understand that the alternative type will add the positive inflation premium to the tag-along authority's choices.

6 Summary and forward-looking statements

We have studied optimal monetary policy in a an imperfect public monitoring framework, where skeptical private agents learn rationally about the nature of the monetary authority and the monetary chooses its actions taking into account private sector learning.[24] Our focus was on issues of imperfect credibility that are plausibly relevant for the 1970s through the early 2000s, in that we examined disinflation dynamics and stabilization policy. A key result was that the optimal pattern of inflation management depended critically on the nature of the skepticism of the private sector, whether it was principally concerned about a mechanically inflationary alternative monetary authority (the automaton of section 4) or about one which would tag-along with the strong decisionmaker's actions. This reinforces our view that an understanding of optimal policy under imperfect credibility requires an analysis of the type of strategic interaction between types of policy authorities that we have begun to examine in companion research.

Recent events in advanced economies have placed alternative challenges in front

[24]We adopted the imperfect public monitoring framework for three reasons. First, it seems to us to capture actual aspects of the real world, where the private sector does not know the exact actions taken by a monetary authority, even one that operates under an interest rate rule because policy is multidimensional even in that setting. Second, in an imperfect public monitoring setting, there is the potential for reputation to rise and fall as a result of shocks even without changes in underlying policymaker behavior. Third, it complements other work on investment in reputation that we have done with other assumptions, including the use of mixed strategies in Lu [2012] and the use of stochastic heterogeneity of the alternative type (King, Lu, and Pasten [2008]).

of the world's central banks, in both the conduct of monetary and banking policy. Notably, the difficulties of conducting monetary and banking policy at the zero lower bound as well as the ongoing challenges to the European monetary system are evidently very different from the problems of the 1980s. Yet, we see issues of imperfect credibility as central to each of these more recent developments, so that these also motivate our inquiries into the design of optimal policy in settings with of private sector skepticism.

Tables

Parameter	First Equation	Definition	Benchmark
β	(1)	Discount factor	.995
h	(2)	Output weight	.004
x^*	(2)	Output target	.1
κ	(3)	PC output slope	.04
δ	Appendix B	Persistence of inflation shock	.9
σ_ς		Std of inflation shock	.02
σ_ε		Std of implementation error	.01

Variable	First Equation	Definition
π	(1)	Inflation
x	(1)	Output gap
u	(1)	Momentary objective
e	Appendix A	Expected inflation
τ	(4)	Type of the monetary authority
ς	(3)	Inflation shock
a	(4)	Policy action
a_1, a_2	(5)	Policy action of the strong type and the alternative type
ε	(4)	Implementation error
ρ	(5)	Credibility
s	(5)	True state of the economy
γ	(8)	Commitment multiplier
η	(8)	Pseudo-state variable
Ω	(8)	Intermediate value function
W	(8)	Value function
m	(8)	Short-hand for $E\pi_{t+1}\|a_2$
F	(10)	cdf for implementation error
b	(6)	Bayesian updating function

Table 2.1

Coefficients for the multiplier γ_t

$$\underline{\gamma} = \frac{q\kappa x^* + q\beta(1-\rho)\mu}{1 + q\,(1-l)\,(1-\beta l)}$$

$$\theta_{\gamma\eta} : \text{the stable root of } \beta ql \ z^2 - \left[\beta ql^2 + (1+q)\right] z + ql = 0$$

$$\theta_{\gamma\varsigma} = \frac{\beta q\,(1-\rho)\,\phi\delta + q}{1 + q - q\beta l\,(\theta_{a\eta} + \delta)}$$

$$\theta_{\gamma\varepsilon} = \frac{-q}{1 - \beta ql\theta_{a\eta}}$$

Coefficients for the policy action a_t

$$\underline{a} = (1-l)\,\underline{\gamma}$$

$$\theta_{a\eta} = \theta_{\gamma\eta} - l$$

$$\theta_{a\varsigma} = \theta_{\gamma\varsigma}$$

Coefficient for the output gap x_t

$$\underline{x} = \frac{1}{\kappa}[\underline{a} - \beta l\underline{a} - \beta(1-\rho)\mu]$$

$$\theta_{x\eta} = \frac{1}{\kappa}[\theta_{a\eta} - \beta l\theta_{a\eta}\theta_{\gamma\eta}]$$

$$\theta_{x\varsigma} = \frac{1}{\kappa}[\theta_{a\varsigma} - \beta l\theta_{a\eta}\theta_{\gamma\varsigma} - \theta_{a\varsigma}\delta - 1]$$

$$\theta_{x\varepsilon} = \frac{1}{\kappa}(1 - \beta l\theta_{a\eta}\theta_{\gamma\varepsilon})$$

Coefficient for the discretionary rule

$$\mu = \frac{\frac{\kappa}{\varepsilon}}{(1 - \beta + \frac{\kappa^2\varepsilon}{\lambda})} = \frac{\kappa h x^*}{\kappa^2 + (1-\beta)\,h}$$

$$\phi = \frac{1}{\varepsilon\kappa + 1 - \beta\delta} = \frac{h}{\kappa^2 + (1-\beta\delta)\,h}$$

Table **3.1**

	$\rho=0$	$\rho=0.5$	$\rho=1$
$\underline{a}$	0.99%	0.68%	0
$\theta_{\gamma\eta}$	0	0.34	0.54
$\theta_{a\varsigma}$	1.98	1.83	1.04
$\theta_{a\eta}\theta_{\gamma\varsigma}-\theta_{\gamma\eta}\theta_{a\varsigma}$	0	-0.91	-1.04

	$\rho=0$	$\rho=0.5$	$\rho=1$		$\rho=0$	$\rho=0.5$	$\rho=1$
$\underline{\gamma}$	0.0099	0.0137	0.01	$\underline{x}$	0.0012	-0.0369	0
$\theta_{\gamma\eta}$	0	0.34	0.54	$\theta_{x\eta}$	0	-**3.3**8	-**5.3**8
$\theta_{\gamma\varsigma}$	1.98	1.83	1.04	$\theta_{x\varsigma}$	-19.82	-18.27	-10.37
$\theta_{\gamma\varepsilon}$	-2.5	-2.08	-1.16	$\theta_{x\varepsilon}$	25	20.8	11.63

Table **3.2**

Case	Figure		ω_p	ω_b
automaton	**4.2**		0	0
no learning effect of policy actions	**5.1**	a does not affect b	0	1
no leverage loss from imperfect credibility	**5.2**	$l=1$	1	0
full tag-along behavior	**5.3**		1	1

Table **5.1**

7 References

Abreu, Dilip, David Pearce and Ennio Stacchetti, 1990. "Toward of Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58: pp. 1041-1063.

Atkeson, Andrew and Patrick J. Kehoe, 2006. "The Advantage of Transparency in Monetary Policy Instruments," The Federal Reserve Bank of Minneapolis Staff Report 297.

Backus, David A, and John Driffill. 1985a. "In‡ation and Reputation." *American Economic Review*, 75(3): 530-538.

Backus, David A, and John Driffill. 1985b. "Rational Expectations and Policy Credibility Following a Change in Regime." *The Review of Economic Studies*, 52(2): 211-221.

Ball, Laurence. 1994. "Credible Disinflation with Staggered Price-Setting." *American Economic Review*, 84(1): 282-289.

Ball, Laurence. 1995. "Disinflation with imperfect credibility." *Journal of Monetary Economics*, 35(1): 5-23.

Barro, Robert J. 1985. "Reputation in a Model of Monetary Policy with Incomplete Information." *Journal of Monetary Economics*, 17: 3-20.

Baxter, Marianne, 1985 "The role of expectations in stabilization policy," *Journal of Monetary Economics*, 15(3): 343-362.

Bernanke, Ben S. 2004 "Gradualism", http://www.federalreserve.gov/boarddocs/speeches/2004/20

Chari, V.V., Kehoe, Patrick.J., 1990. "Sustainable plans". *Journal of Political Economy* 98, 783–802.

Clarida, Richard, Jordi Gali and Mark Gertler, 1999. "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, vol. 37(4), pages 1661-1707,

Cogley, T., Matthes, C. and A.M. Sbordone, 2011. "Optimal Disinflation Under Learning," mimeo

Cripps, Martin W., George J. Mailath and Larry Samuelson. 2004. "Imperfect Monitoring and Impermanent Reputations." *Econometrica*, 72(2): 407-432.

Cukierman, Alex, and Allen H. Meltzer. 1986. "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 54(5): 1099-1128.

Cukierman, Alex, and Nissan Liviatan. 1991. "Optimal Accommodation by Strong Policymakers under Incomplete Information." *Journal of Monetary Economics*, 27(1): 99-127.

Fudenberg, Drew, David Levine and Eric Maskin, 1994. "The Folk Theorem with Imperfect Public Information," *Econometrica* 62: pp. 997—1039.

Gali, Jordi, (2008), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.

Giannoni, Marc and Michael Woodford. 2004. "Optimal Inflation-Targeting

Rules." NBER Chapter, in The Inflation-Targeting Debate: 93-172..

Khan, A., King, R. G. and Wolman, A. L. (2003), "Optimal Monetary Policy", *The Review of Economic Studies*, 70: 825–860. doi: 10.1111/1467-937X.00269.

King, R. G., Lu, Y. K. and Pasten, E. S. (2008), "Managing Expectations", *Journal of Money, Credit and Banking*, 40: 1625–1666.

King, Robert G., Yang K. Lu, and Ernesto S. Pasten, (2011) "Optimal policy design with a skeptical forward-looking private sector," partial manuscript June.

King, Robert G. and Alexander L. Wolman, (1999) "What Should Monetary Policy Do If Prices Are Sticky?"(with Alexander L. Wolman), John B. Taylor, ed., *Monetary Policy Rules*, University of Chicago Press for National Bureau of Economic Research, 349-404.

Kurozumi, Takushi, 2008. "Optimal sustainable monetary policy," *Journal of Monetary Economics* 55(7): pp. 1277-1289, October.

Loisel, Olivier, 2008. "Central bank reputation in a forward-looking model," *Journal of Economic Dynamics and Control* 32(11):pp. 3718-3742, November.

Lu, Yang K., 2012. "Optimal Policy with Credibility Concerns," Hong Kong University of Science and Technology, Department of Economics, working paper.

Mailath, George J. and Larry Samuelson, 2001. "Who Wants a Good Reputation?" *The Review of Economic Studies*, 68(2): 415-441.

Milani, Fabio, 2007. "Expectations, learning and macroeconomic persistence," *Journal of Monetary Economics*, vol. 54(7), pages 2065-2082, October.

Phelan, Chirstopher. 2005. "Public trust and Government Betrayal." *Journal of Economic Theory*, 130: 27-43.

Sargent, Thomas J., 1982, "The Ends of Four Big Inflations," in *Inflation: Causes and Consequences*, ed. by Robert E. Hall (Chicago: University of Chicago Press), pp. 41–97.

Schaumberg and Tambalotti (2007) "An Investigation of the gains from commitment,"

Schorfheide, Frank, 2005. "Learning and Monetary Policy Shifts," *Review of Economic Dynamics*, Elsevier for the Society for Economic Dynamics, vol. 8(2), pages 392-419, April.

Svensson, L. E. O., 1997, "Optimal inflation targets, "conservative" central banks, and linear inflation contracts," *American Economic Review* 87, 98–114.

Walsh, Carl, *Monetary Theory and Policy*, 2nd Edition, MIT Press, 2003

Woodford, Michael, (1999), “Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?”in Federal Reserve Bank of Kansas City (ed), *New Challenges for Monetary Policy*, Kansas City.

Woodford, M. (2003a), "Optimal Interest-Rate Smoothing", Review of Economic Studies, 70: 861–885. doi: 10.1111/1467-937X.00270

Woodford, M. (2003b), *Interest and Prices* Princeton University Press.

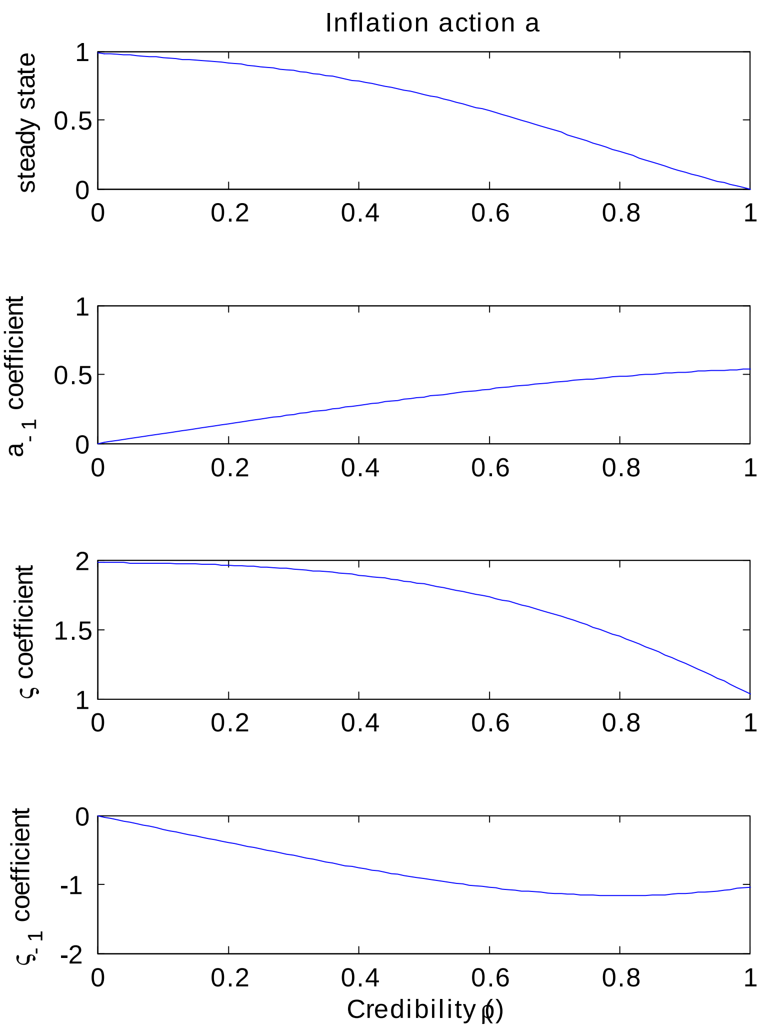

Figure 3.1: Coefficients in the policy rule with fixed action alternative policymaker and exogenous, constant reputation. The top panel is the inflation bias (percent per quarter). The second panel is the response coefficient to the lagged policy action. The third panel is the response coefficient to the momentary inflation shock. The bottom panel is the response coefficient to the lagged inflation shock.

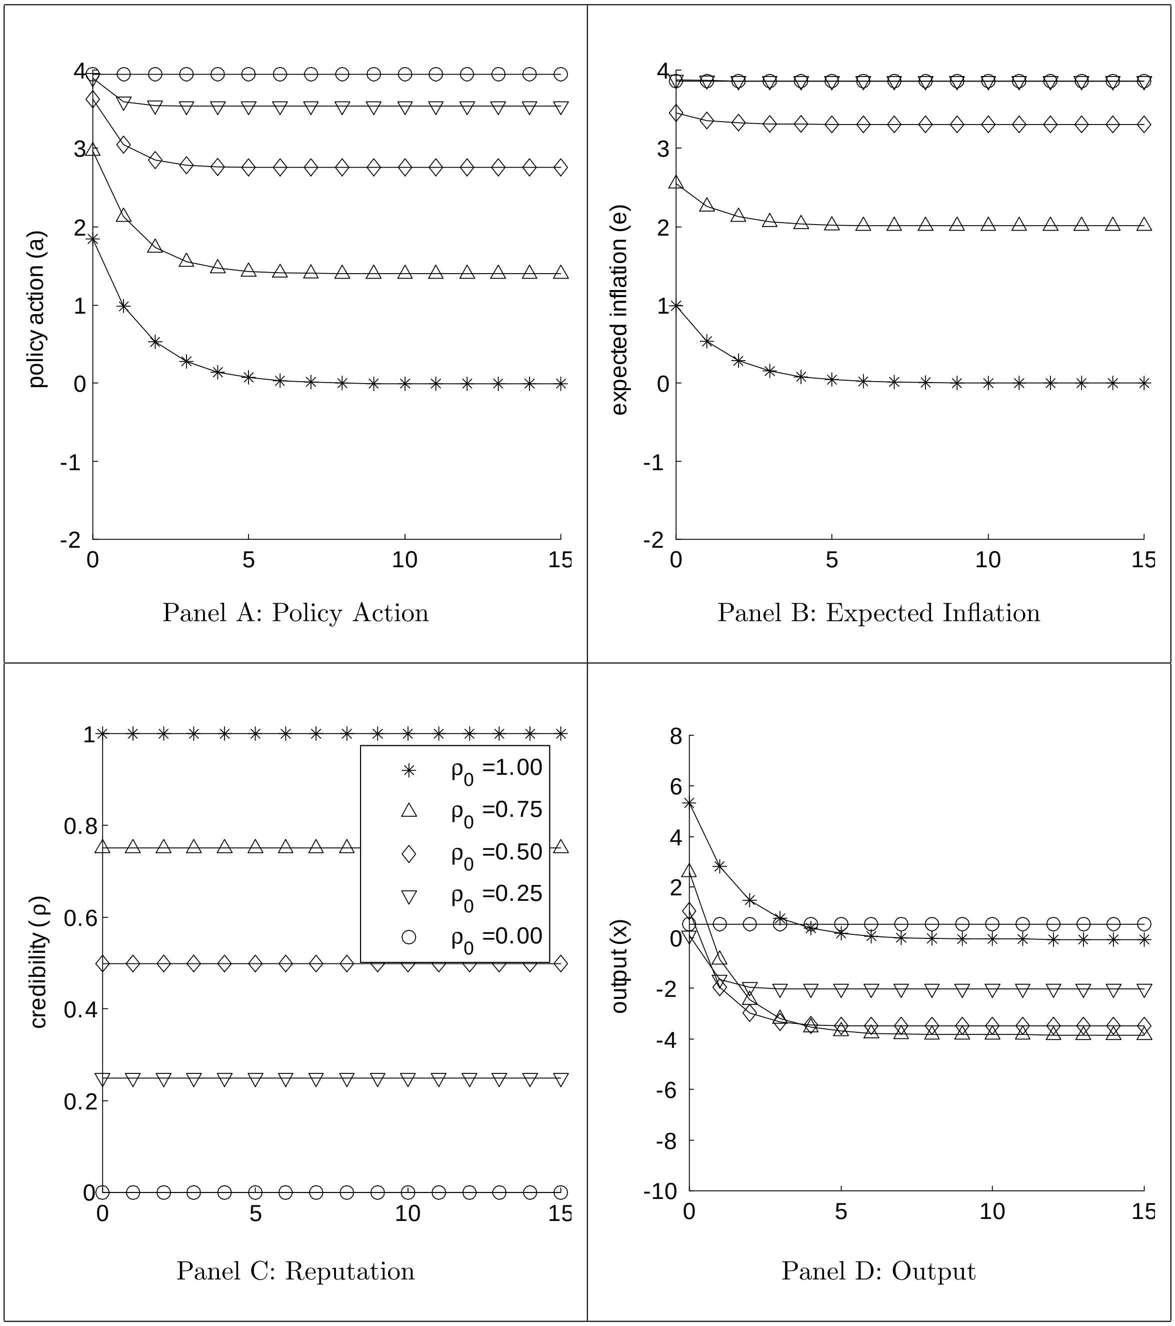

Figure 3.2: Startup dynamics with fixed action alternative policymaker and exogenous, constant reputation. Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

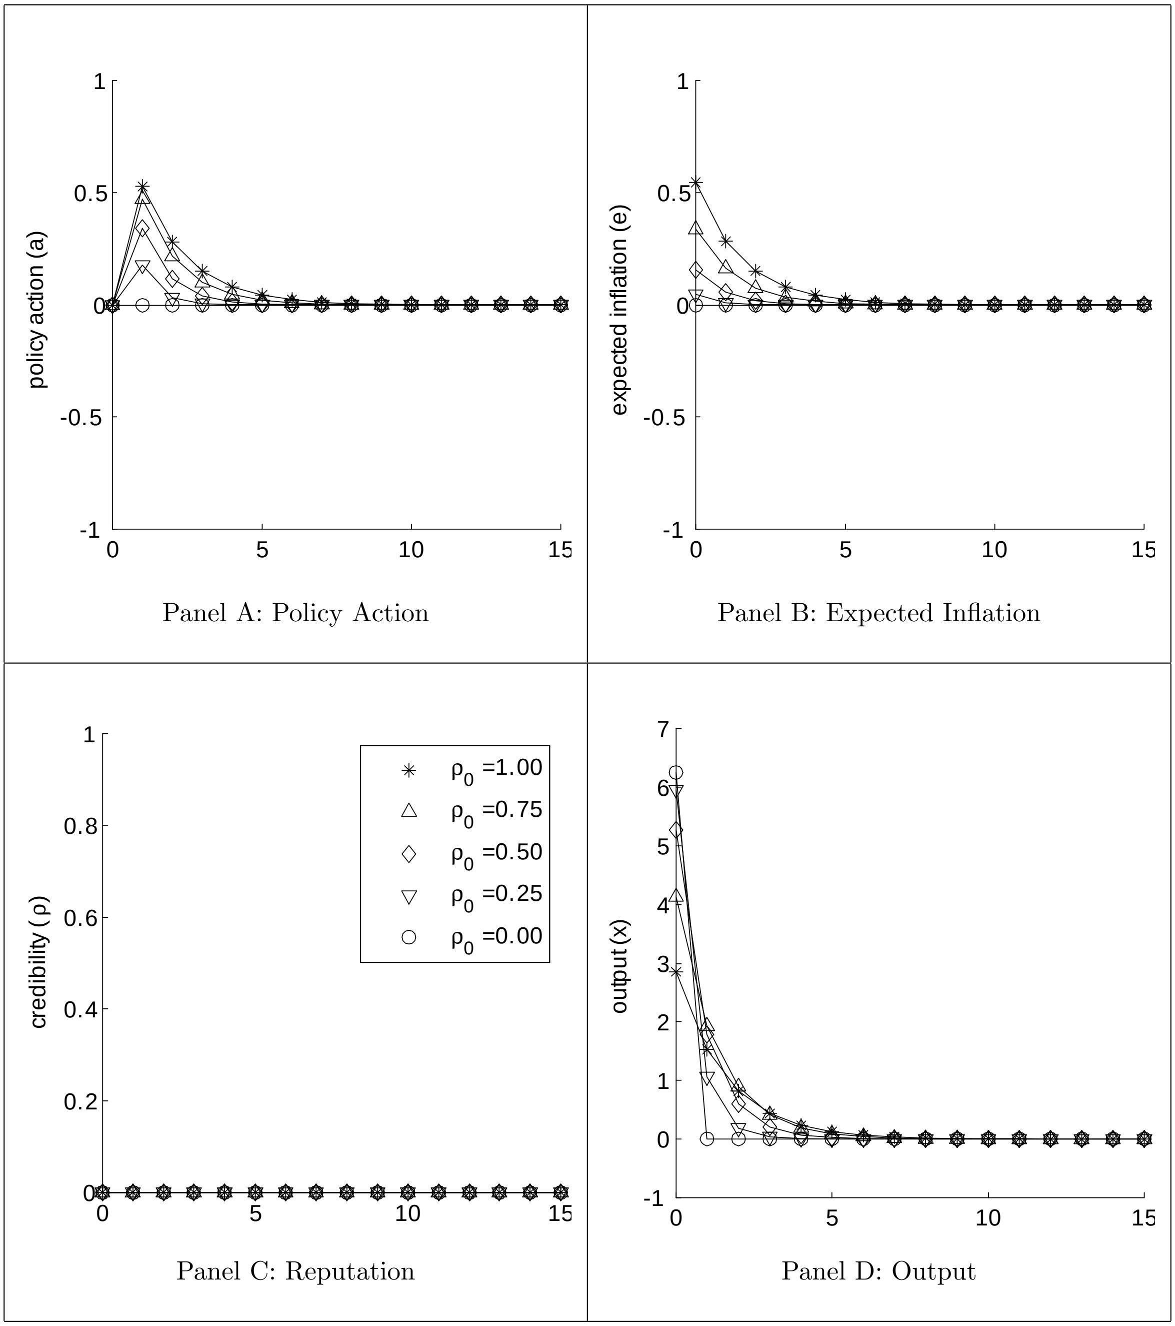

Figure 3.3: Optimal response to positive implementation error with fixed action alternative policymaker and exogenous, constant credibility (error is 1 percent in annual inflation rate or .0025). Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state..

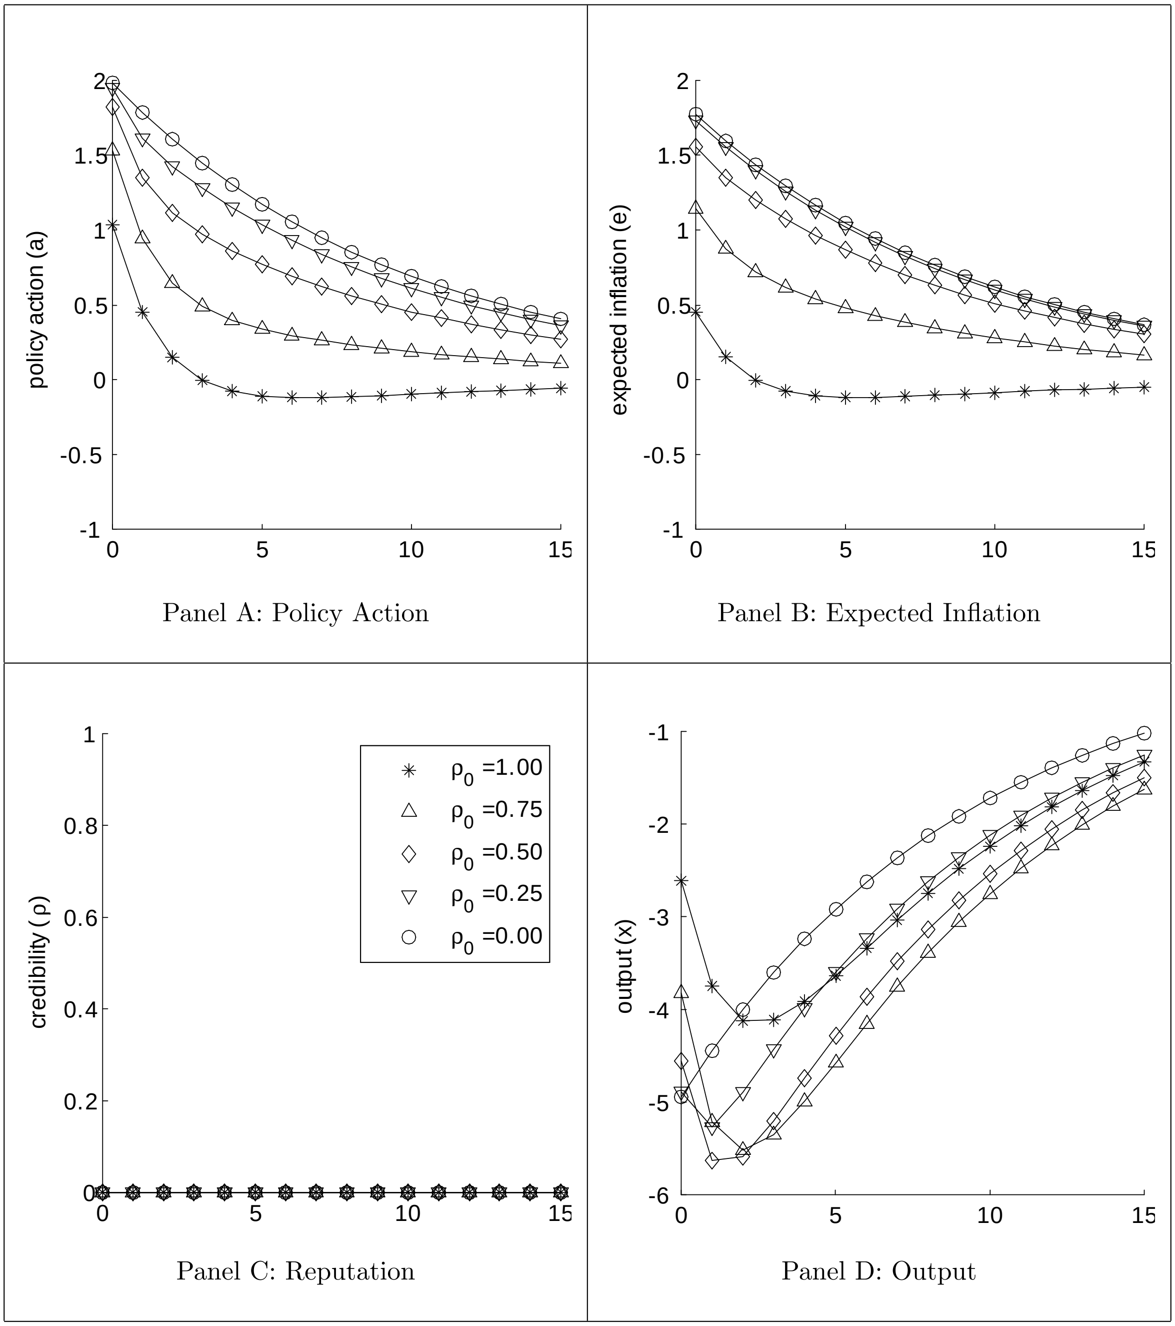

Figure 3.4: Optimal response to positive inflation shock with fixed action alternative policymaker and exogenous, constant credibility (error is 1 percent in annual inflation rate or .0025). Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

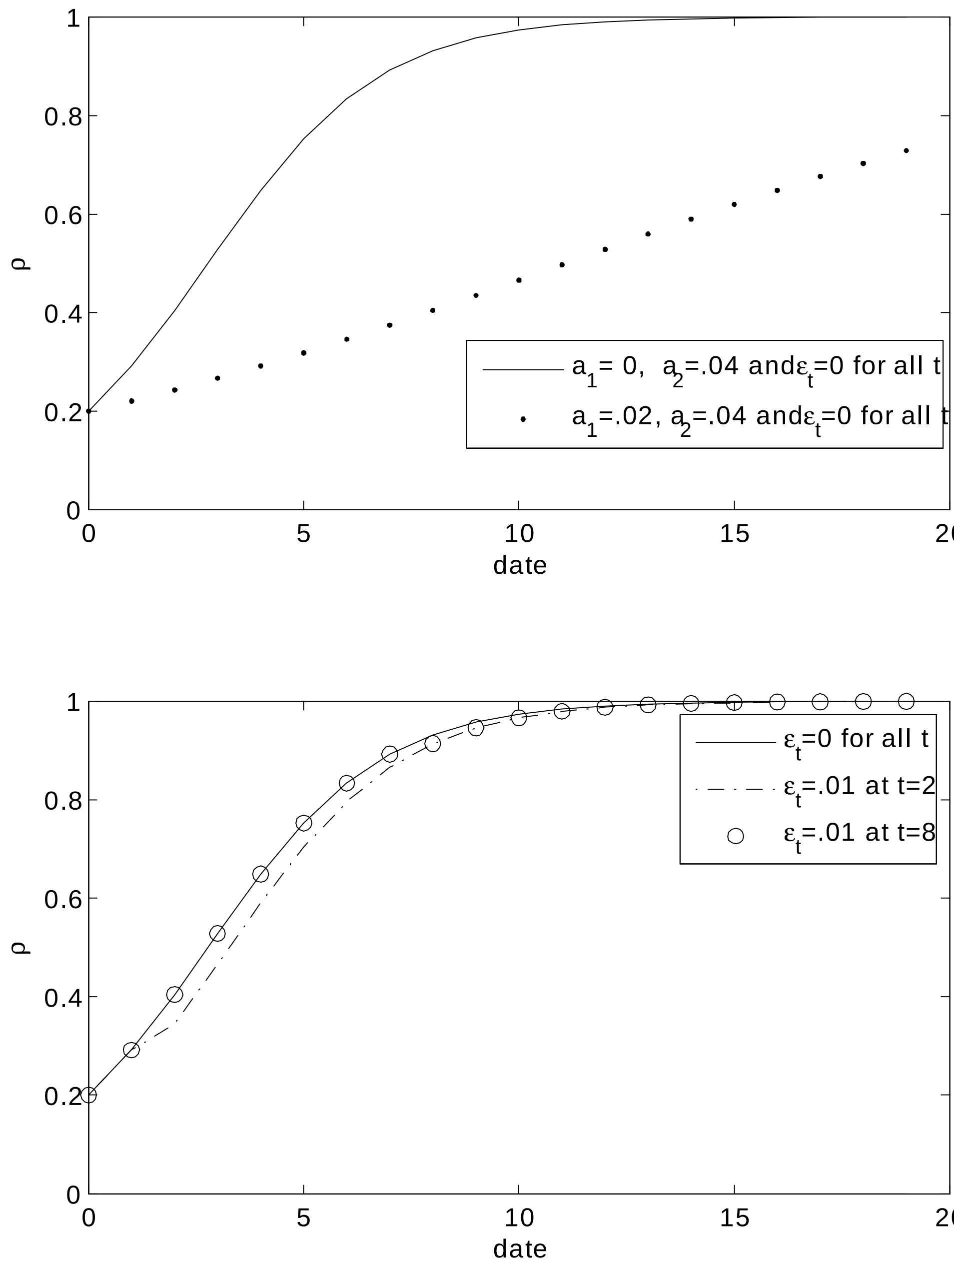

Figure 4.1: Bayesian learning. The top panel shows the alternative learning paths when the strong authority follows either a 2 percent or 0 percent inflation action, assuming that the alternative is setting a 4 percent inflation action. All implementation errors are set at the mean/mode value of 0. The bottom panel shows the effect of an implementation error of 1 percent at two different time points, assuming that the strong authority is using a zero inflation action and the alternative is using a four percent inflation action. Implementation errors in other periods are set to zero. The timing of the model is quarterly, but the inflation rates reported here are annualized (as in other figures and as in the text discussion).

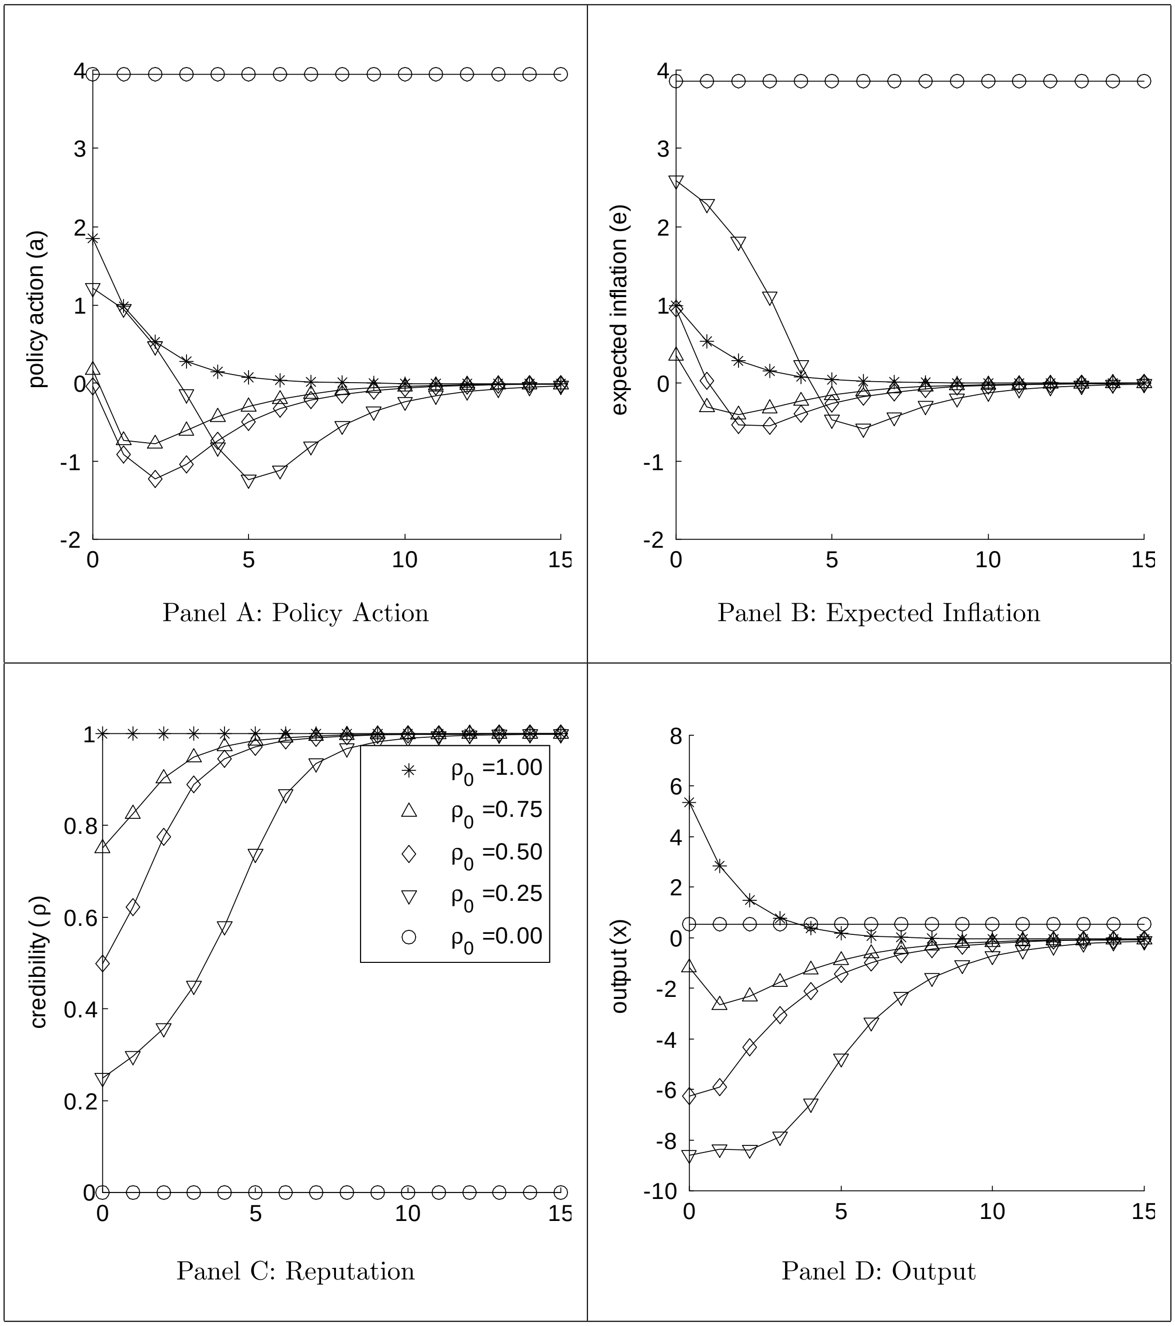

Figure 4.2: Startup dynamics with fixed action alternative policymaker and endogenous credibility. Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

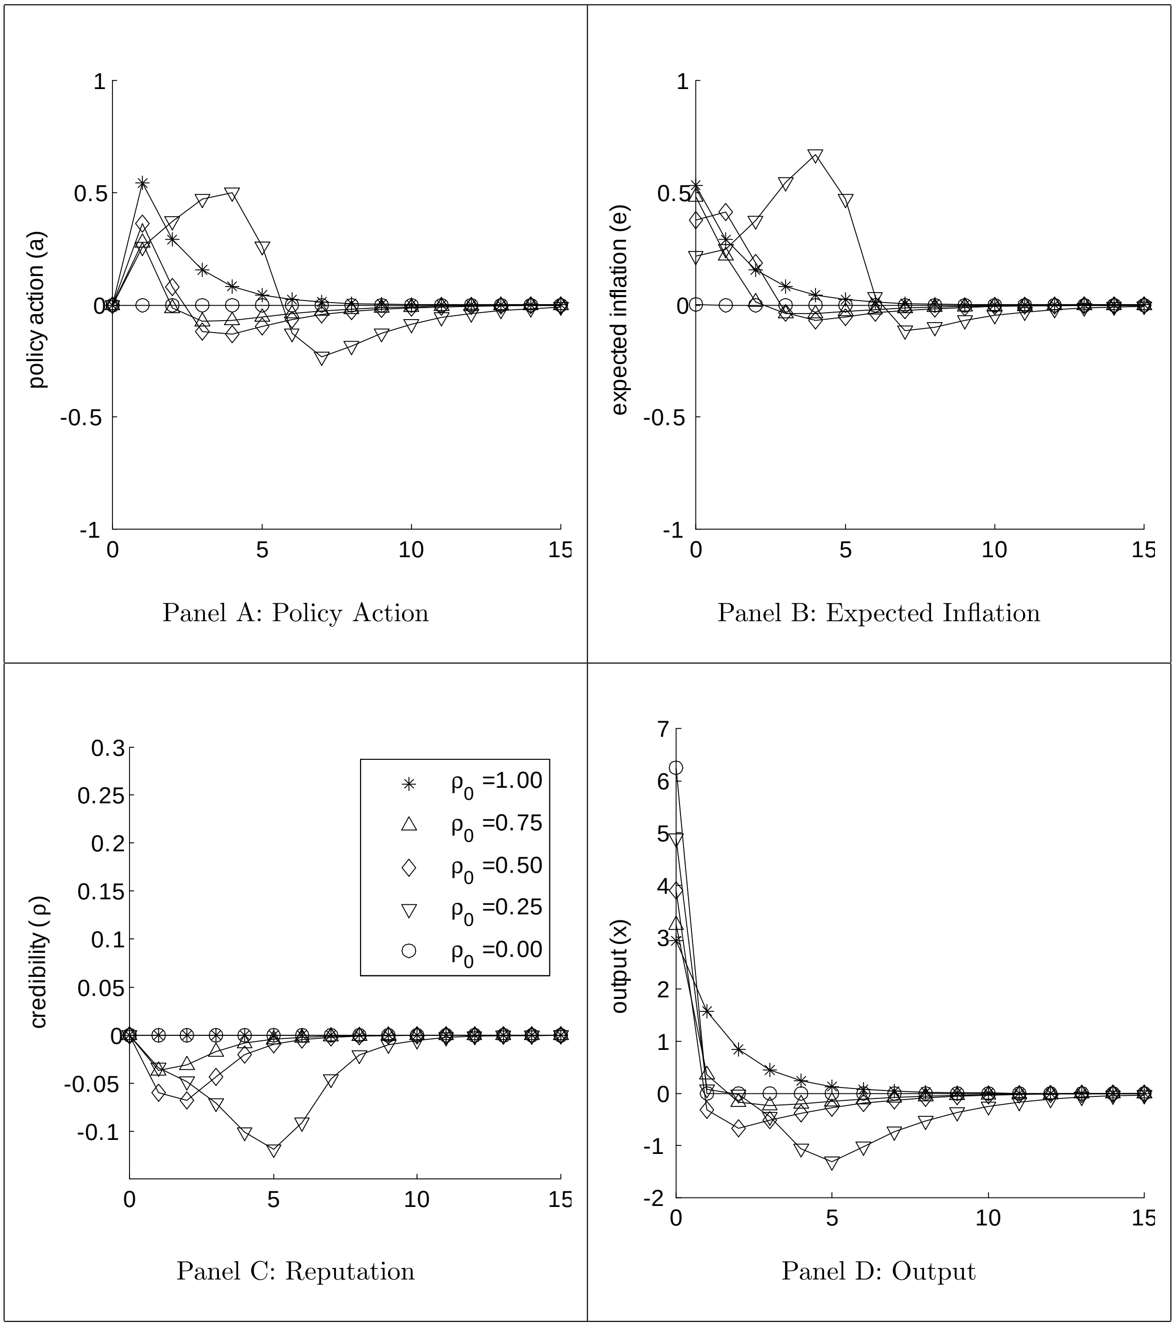

Figure 4.3: Optimal response to positive implementation error with fixed action alternative policymaker and endogenous credibility (error is 1 percent in annual inflation rate or .0025). Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

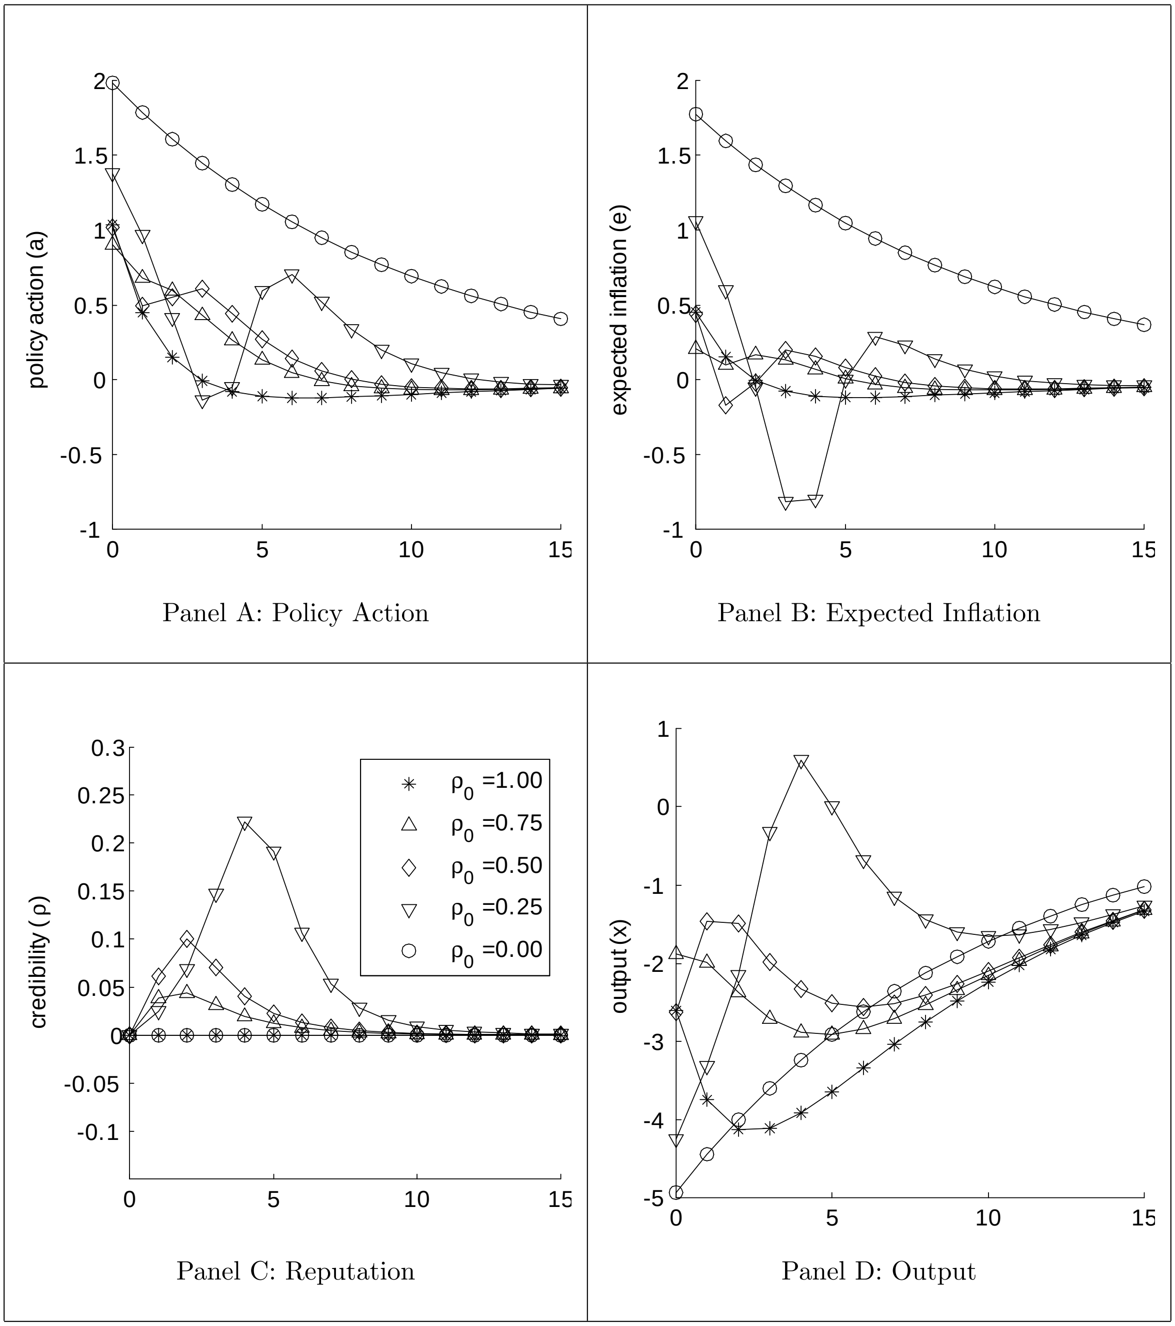

Figure 4.4: Optimal response to positive inflation shock with fixed action alternative policymaker and endogenous credibility (error is 1 percent in annual inflation rate or .0025). Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

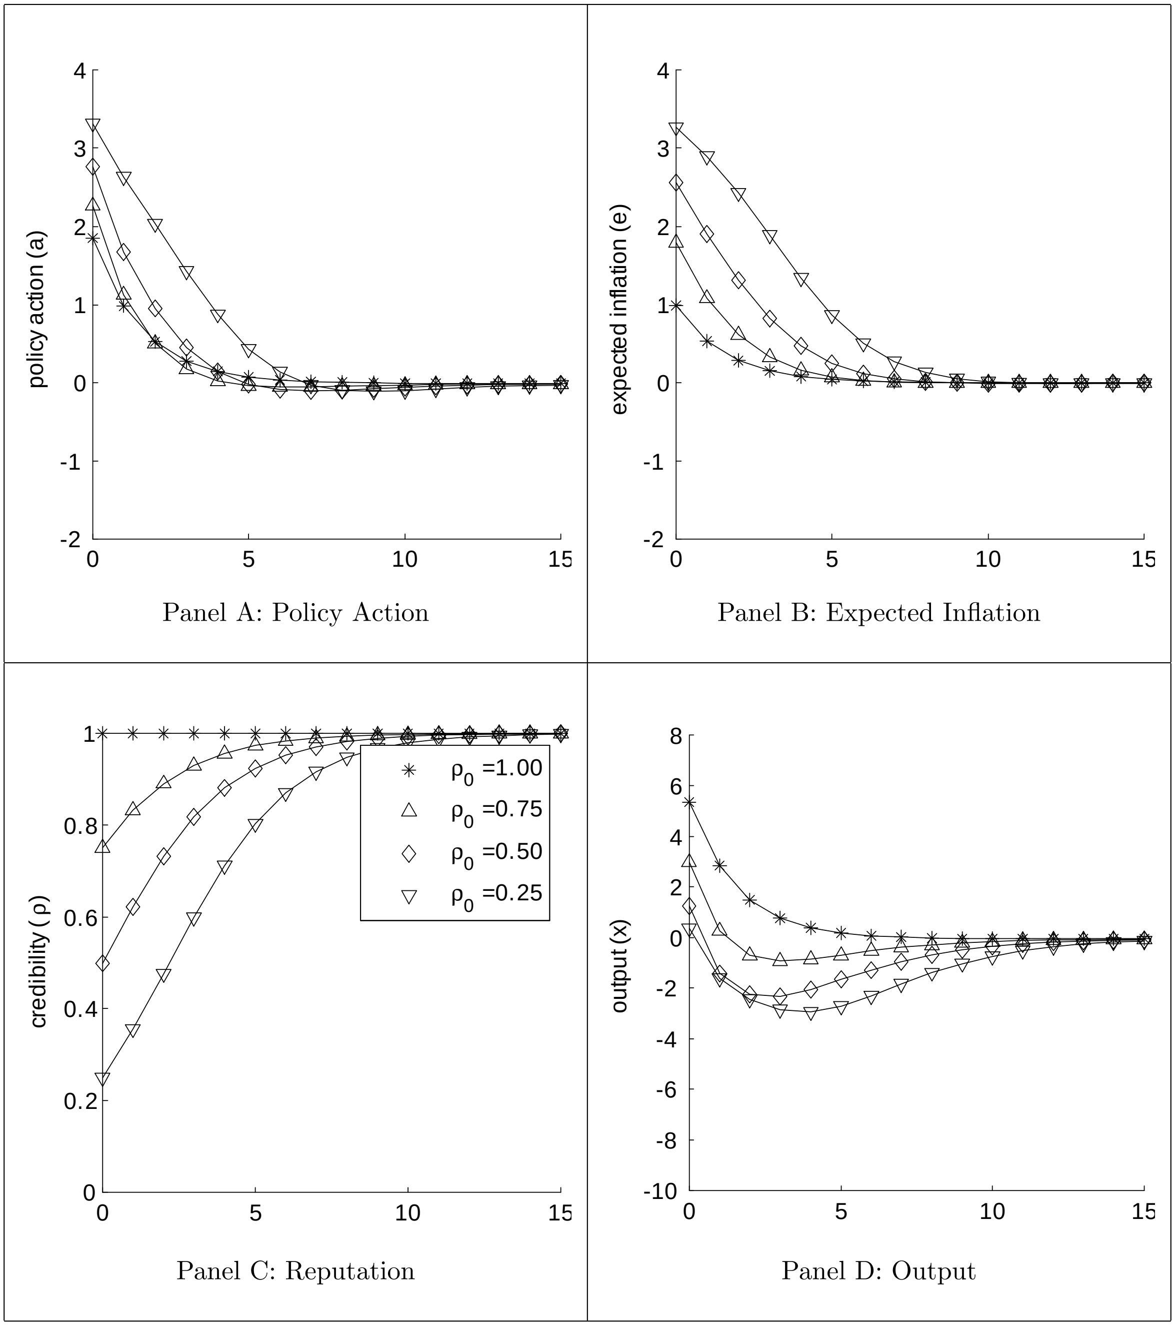

Figure 5.1: Startup dynamics without effect of policy action on learning. Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

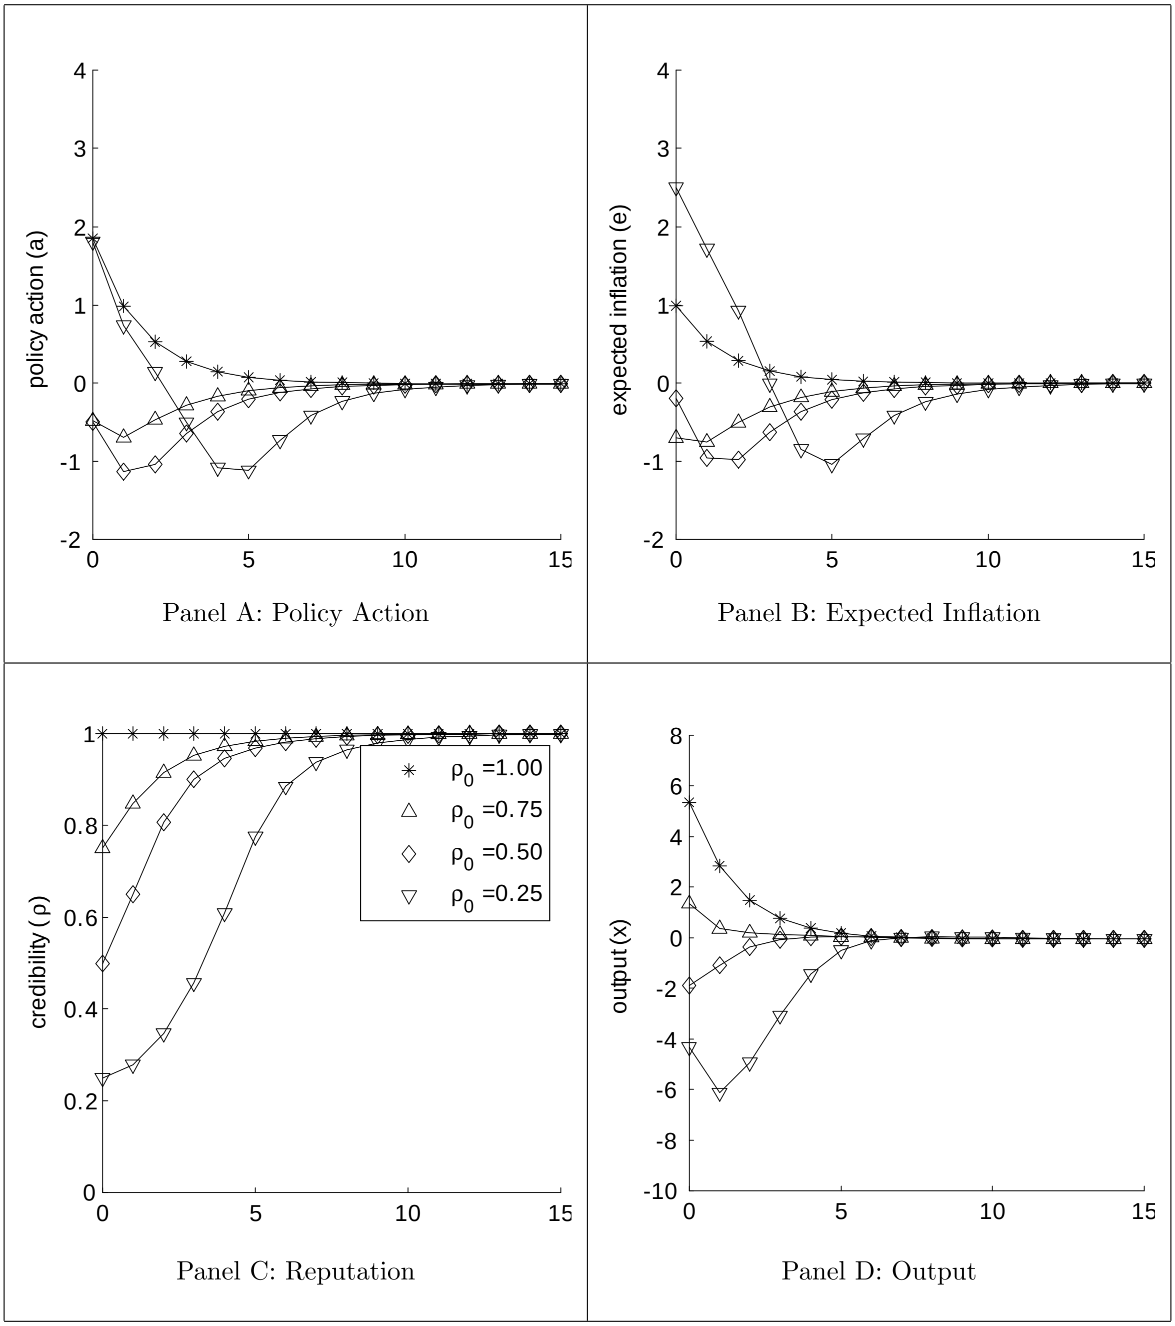

Figure 5.2: Startup dynamics with effect of policy action on learning, but no loss of leverage on expected inflation. Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

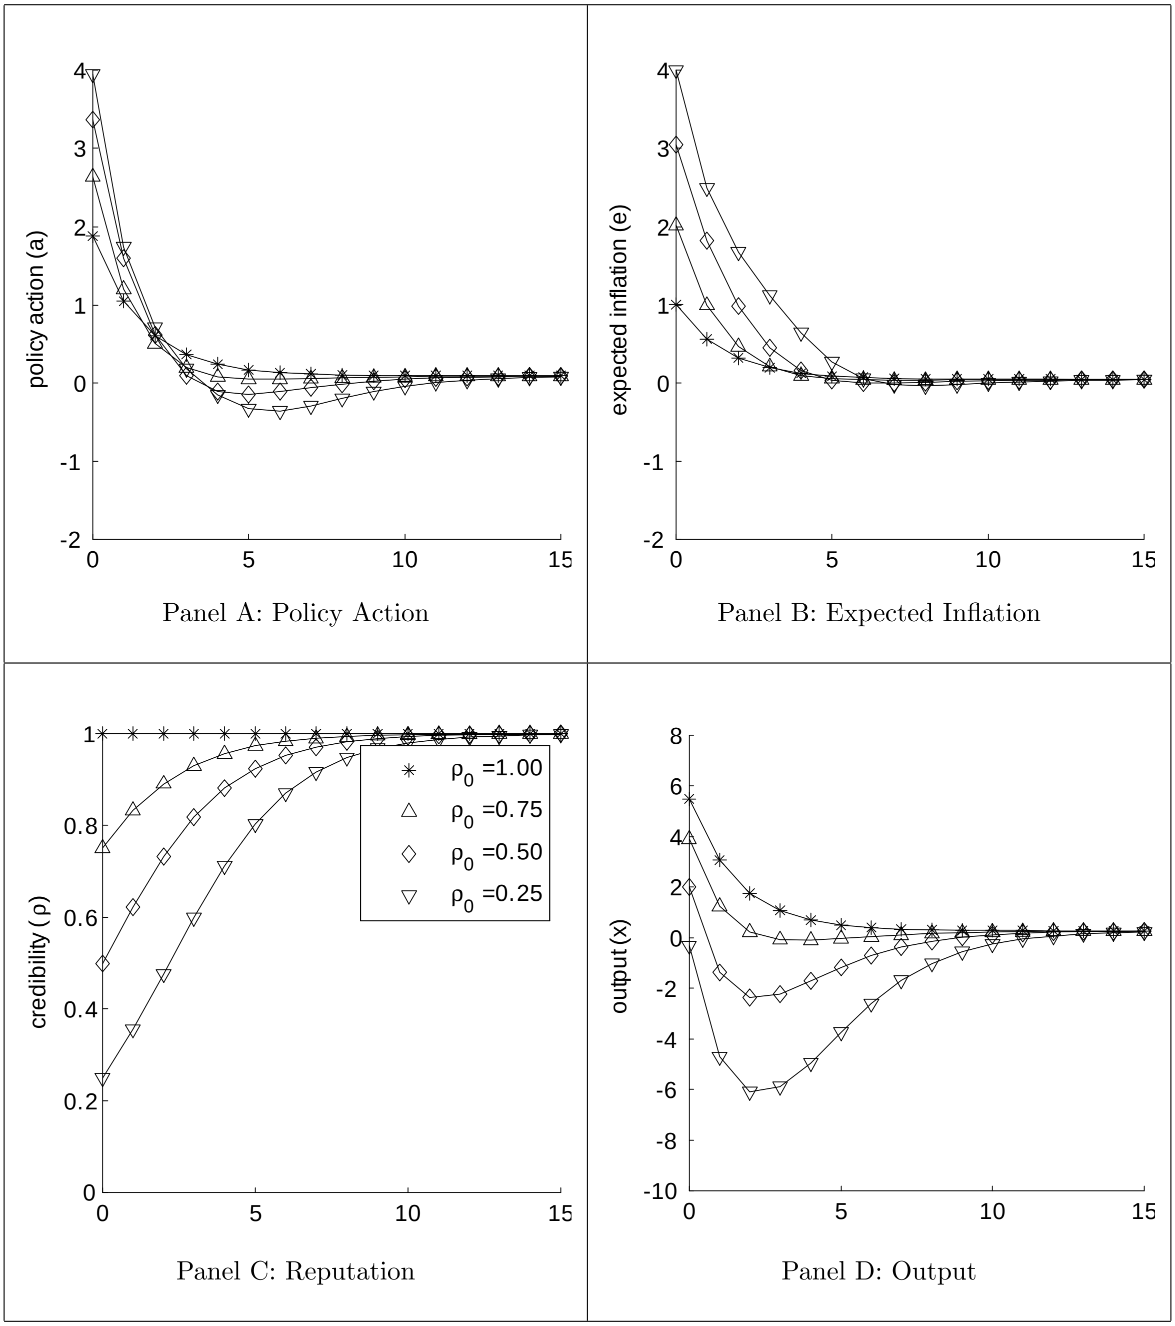

Figure 5.3: Startup dynamics with tag-along alternative policymaker and endogenous credibility. Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

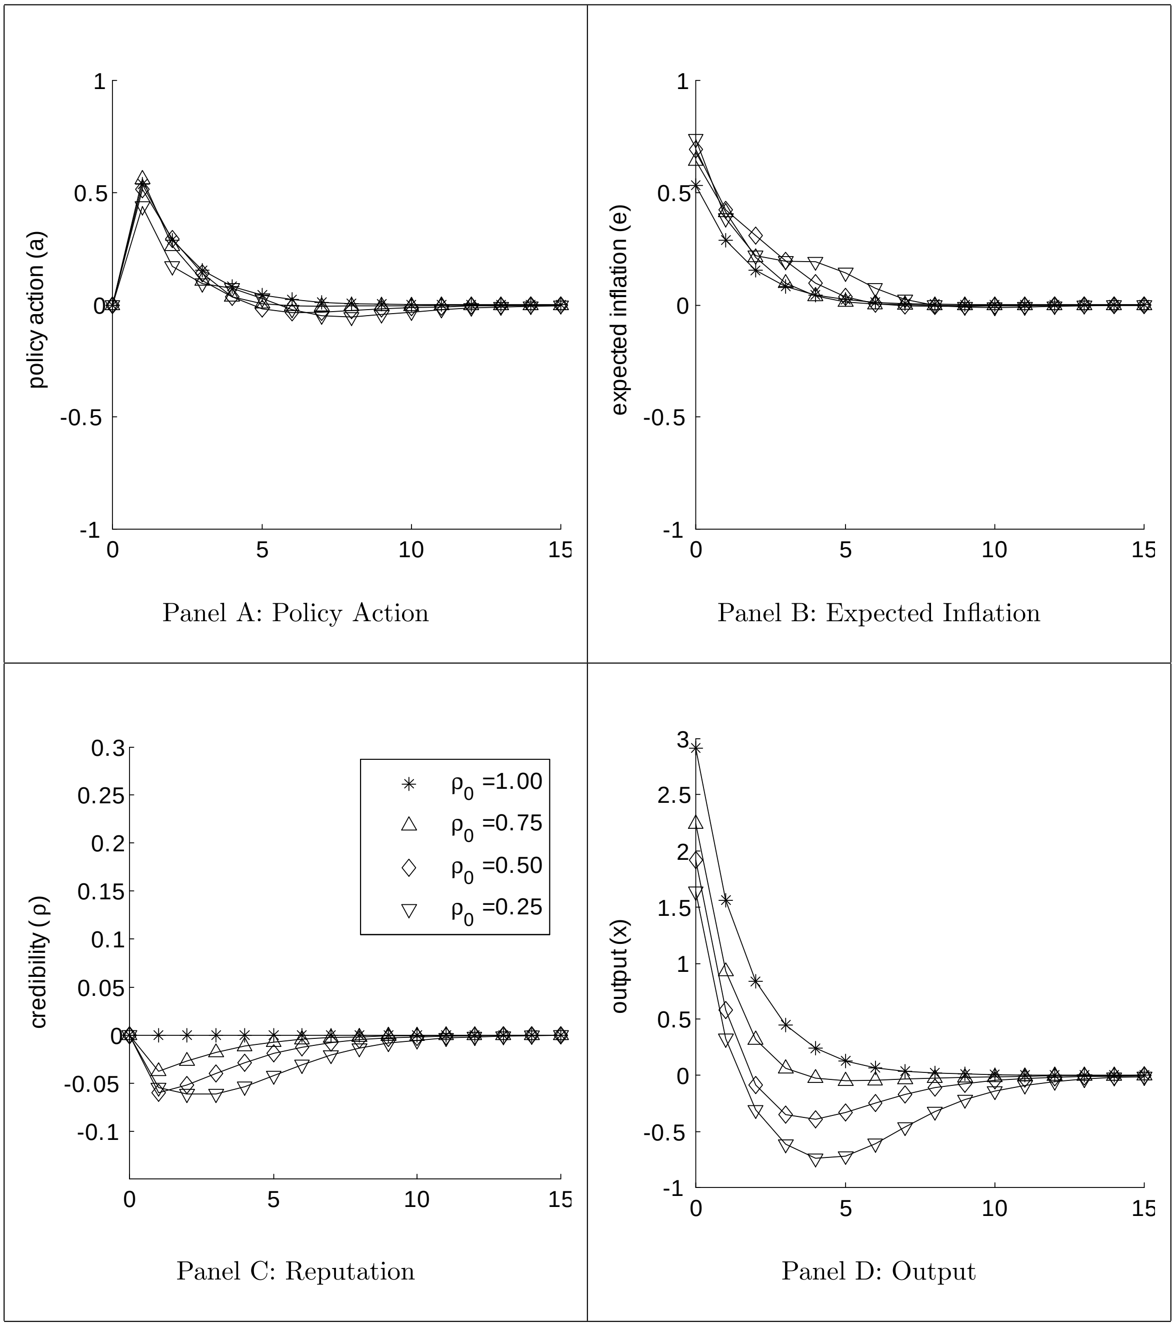

Panel A: Policy Action

Panel B: Expected Inflation

Panel C: Reputation

Panel D: Output

Figure 5.4: Optimal response to positive implementation error with tag-along alternative policymaker and endogenous credibility (error is 1 percent in annual inflation rate or .0025). Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

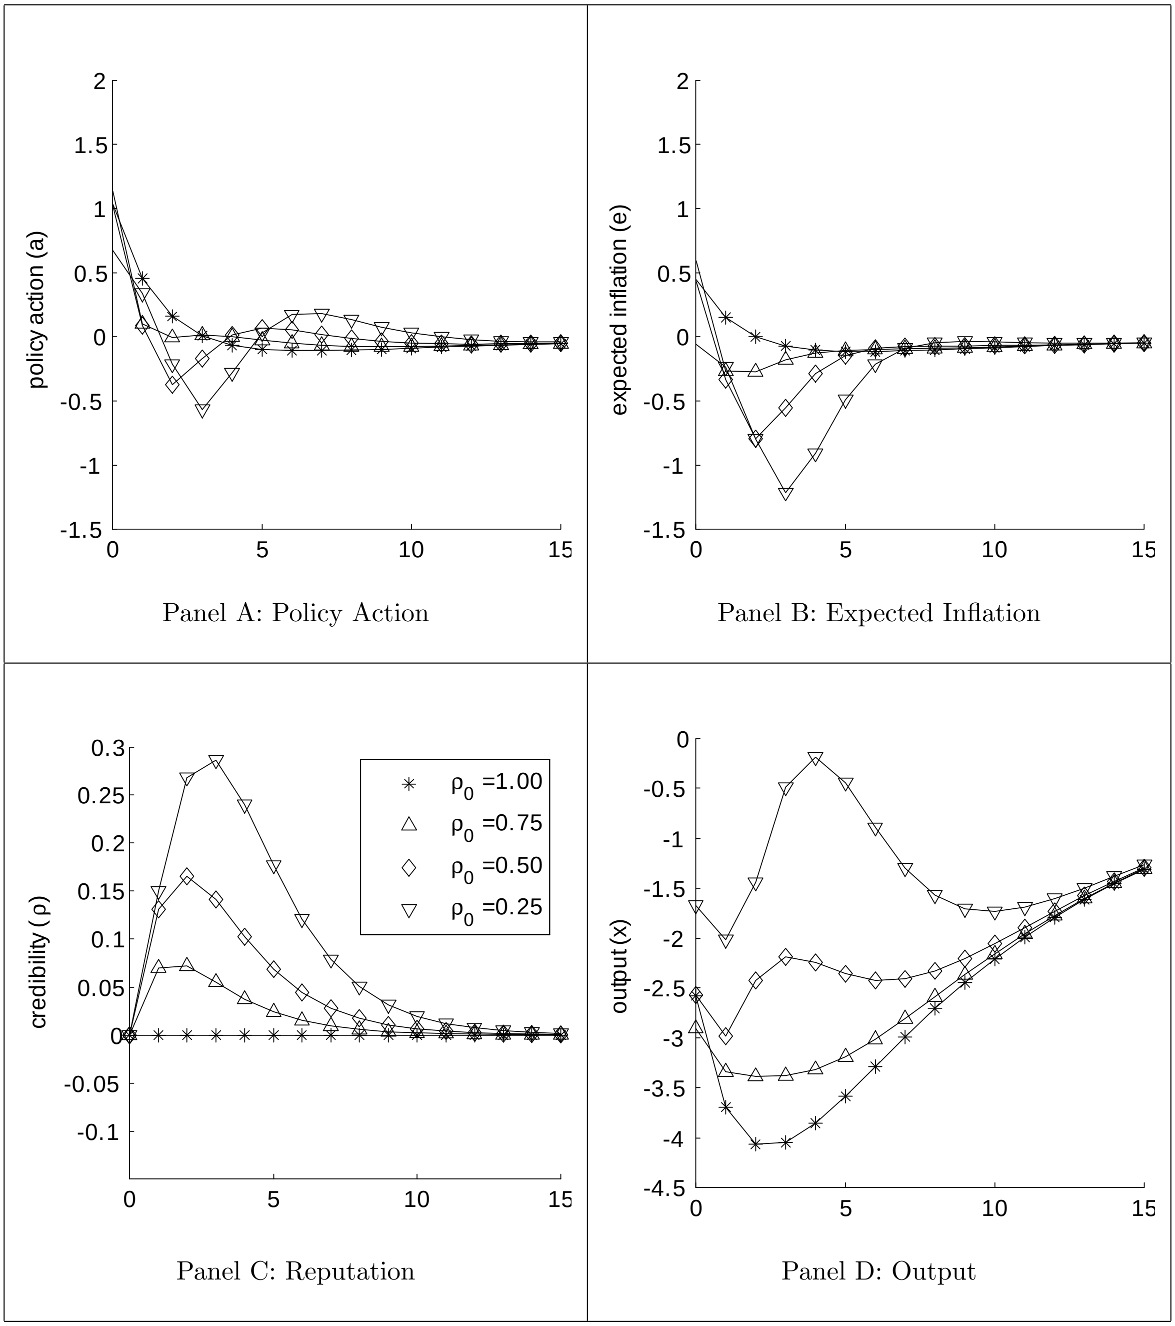

Figure 5.5: Optimal response to positive inflation shock with tag-along alternative policymaker and endogenous credibility (error is 1 percent in annual inflation rate or .0025). Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state..

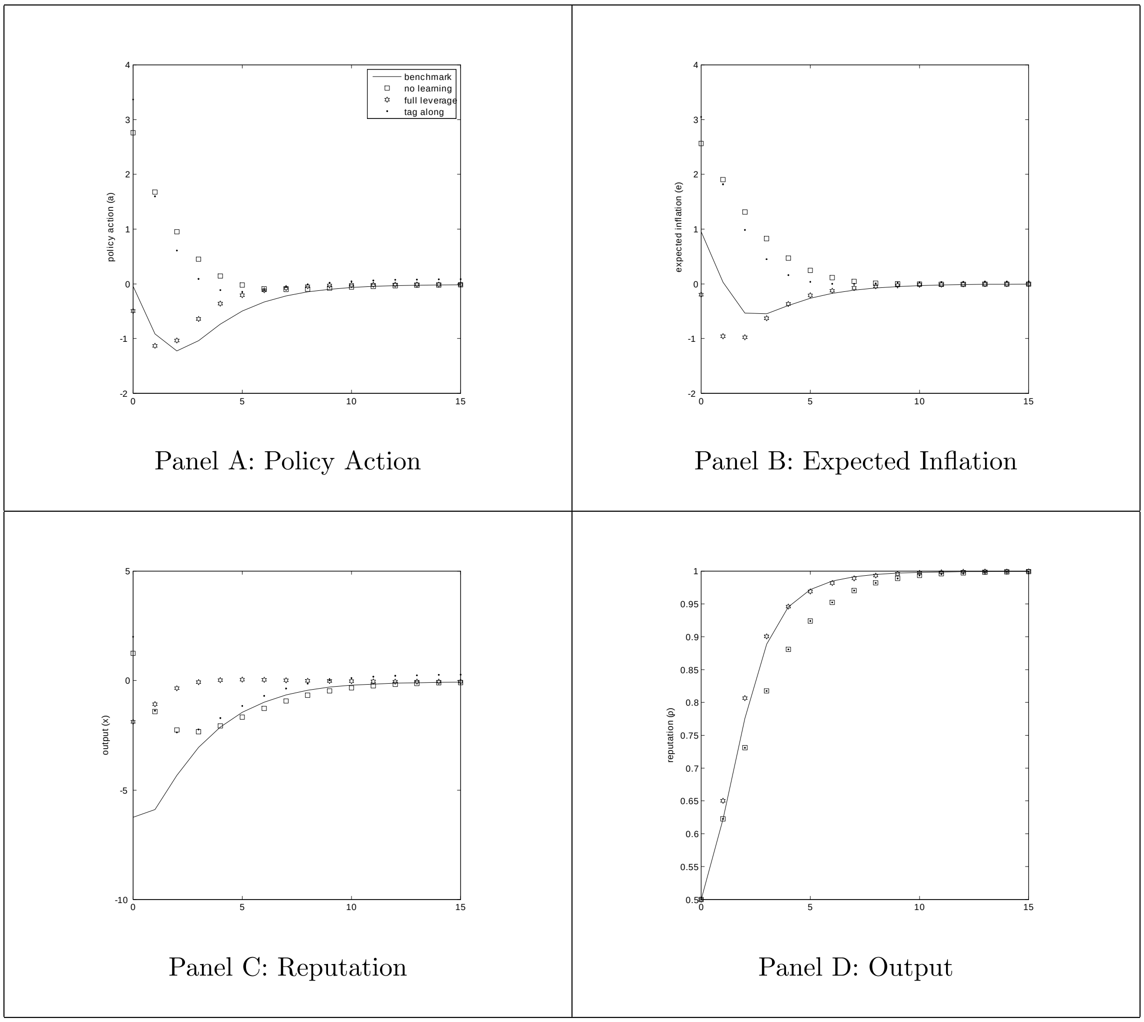

Figure 5.6: Startup dynamics from initial reputation of .5. Four different leverage and learning assumptions previously displayed in Figures 4.1 and 5.1-5.3. Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

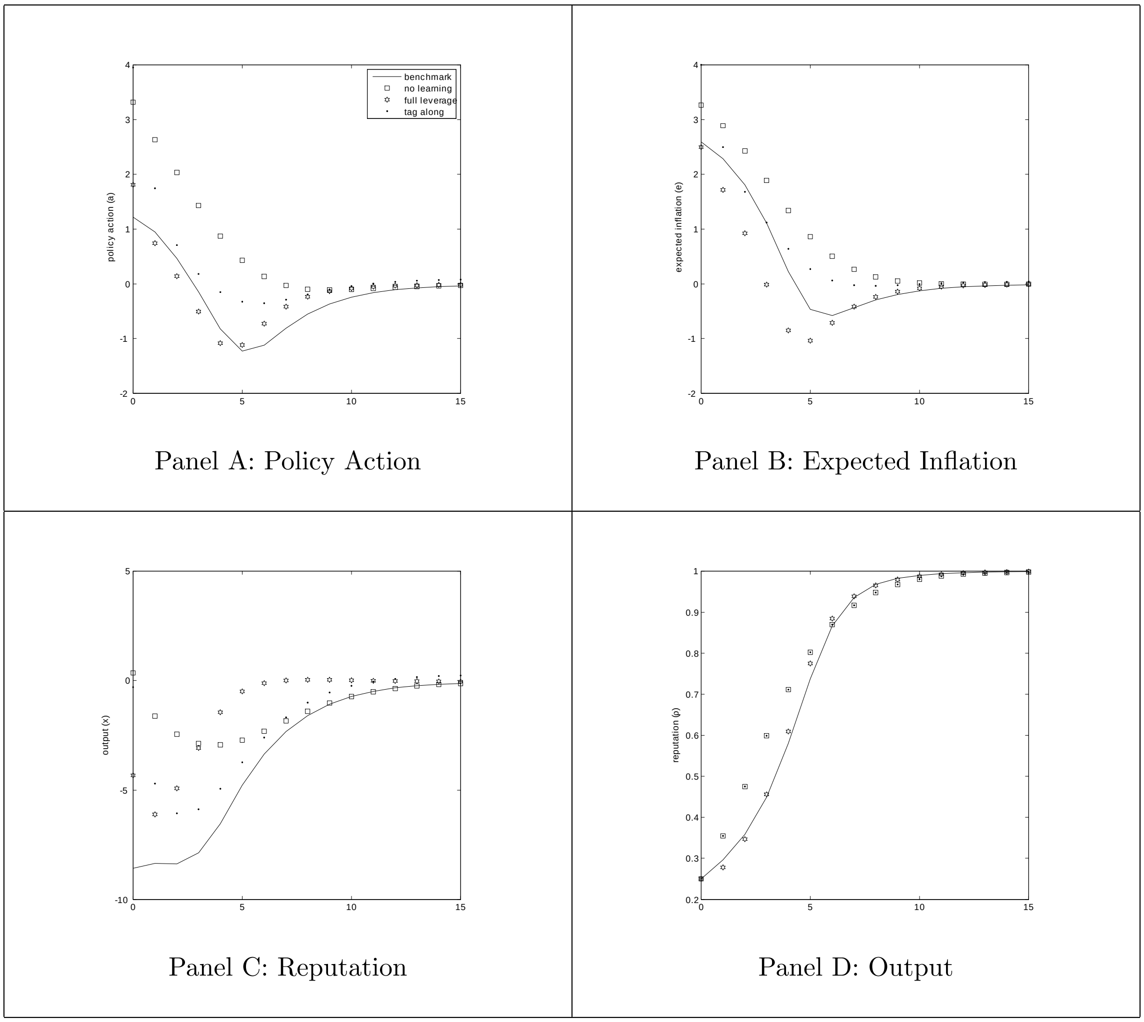

Panel A: Policy Action

Panel B: Expected Inflation

Panel C: Reputation

Panel D: Output

Figure 5.7: Startup dynamics from initial reputation of .25. Four different leverage and learning assumptions previously displayed in Figures 4.1 and 5.1-5.3. Panel A: policy action (mean inflation) is percent per year. Panel B: expected inflation for private agents is percent per year. Panel C: reputation is the likelihood that a trustworthy policymaker is in place. Panel D: output is in percent deviation from distorted steady state.

A Recursive solution method

This Appendix details the derivation of the recursive formation of our optimization problem. We assume that the alternative type of central bank follows some exogenous rule: $a_2(s_t)$ where s_t is the vector of state variables that will be specified in this Appendix.

A.1 Structure of the model

The strong type of central bank is to maximize the expected, present discounted value of a momentary objective

$$\max_{\{\pi_t, x_t\}_{t=0}^{\infty}} E_0\{\sum_{t=0}^{\infty} \beta^t u(\pi_t, x_t)\}$$

with the momentary objective function being:

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

The optimization problem is subject to the standard NK constraint

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + \varsigma_t.$$

Define

$$e_t = \beta E_t \pi_{t+1},$$

x_t can be write as:

$$x_t = \frac{\pi_t - e_t - \varsigma_t}{\kappa}.$$

It turns out to be equivalent and more convenient to express the optimization in term of choosing e_t rather than choosing x_t. The problem thus becomes:

$$\max_{\{\pi_t, e_t\}_{t=0}^{\infty}} E_0\{\sum_{t=0}^{\infty} \beta^t u(\pi_t, e_t, \varsigma_t)\}$$

subject to

$$e_t = \beta \left[\rho_{t+1} E\pi_{t+1}|a_1 + (1-\rho_{t+1}) E\pi_{t+1}|a_2\right] \tag{18}$$

The right hand side of the contraint is determined by the rational expectations of the private sector on the next-period inflation, with the belief about the current central bank being the strong type updated after observing π_t, according to the Baye's rule:

$$\rho_{t+1} = \frac{\rho_t f(\pi_t|a_{1t})}{\rho_t f(\pi_t|a_{1t}) + (1-\rho_t) f(\pi_t|a_{2t})} \tag{19}$$

where $f(\pi, a)$ is the pdf of π conditional on the inflation action being a.

A.2 Recursive method

Denote the public history at the beginning of period t as $h_t = \{h_{t-1}, \pi_{t-1}, \varsigma_t\}$. The timing of the model implies that the inflation action a is a function of h_t and the inflation expectation of the private sector e is a function of h_t and π_t. If we use τ as the type of the central bank: $\tau = 1$ corresponds to the strong type and $\tau = 2$ the weak type,

$$\begin{aligned} a_{1t} &= a_t(h_t, \tau = 1) \\ a_{2t} &= a_t(h_t, \tau = 2) \\ e_t &= e_t(h_t, \pi_t) \end{aligned}$$

Let $p(h_t)$ be the unconditional probability of a history h_t. The expected, present-discounted value of the strong type's momentary objective is

$$E_0 U_0 = \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t) \sum_{\pi_t \in \Pi} f(\pi_t | a_t(h_t, \tau = 1)) \, u(\pi_t, e_t(h_t, \pi_t), \varsigma_t),$$

We attach multipliers $p(h_t) \sum_{\pi_t \in \Pi} f(\pi_t | a_t(h_t, \tau = 1)) \gamma_t(h_t, \pi_t)$ to (18) and produce the Lagrangian component:

$$\begin{aligned} \Psi_0 = & \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t) \sum_{\pi_t \in \Pi} f(\pi_t | a_t(h_t, \tau = 1)) \\ & \gamma_t(h_t, \pi_t) \left[e_t(h_t, \pi_t) - \beta \sum_{\tau} r_t(\tau, h_t, \pi_t) \sum_{\varsigma_{t+1} \in S} \delta(\varsigma_{t+1}, \varsigma_t) \mu(a_{t+1}(h_{t+1}, \tau)) \right] \end{aligned}$$

where

$$\begin{aligned} r_t(\tau = 1, h_t, \pi_t) &= \rho_{t+1} \\ r_t(\tau = 2, h_t, \pi_t) &= 1 - \rho_{t+1} \\ \mu(a) &= E\pi|a = \sum_{\pi_t \in \Pi} f(\pi|a)\,\pi. \end{aligned}$$

Rearranging this, we obtain

$$\begin{aligned} \Psi_0^1 &= -\sum_{t=0}^{\infty} \beta^{t+1} \sum_{h_{t+1}} p(h_{t+1}) \sum_{\tau} \gamma_t(h_t, \pi_t)\, r_t(\tau, h_t, \pi_t)\, \mu(a_{t+1}(h_{t+1}, \tau)) \\ &\quad + \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t) \sum_{\pi_t \in \Pi} f(\pi_t|a_t(h_t, \tau = 1))\, \gamma_t(h_t, \pi_t)\, e_t(h_t, \pi_t) \\ &= E_0 \sum_{t=0}^{\infty} \beta^t \left[-\gamma_{t-1}[\rho_t \mu(a_{1t}) + (1-\rho_t)\mu(a_{2t})] + \sum_{\pi_t \in \Pi} f(\pi_t|a_{1t})\, \gamma_t(\pi_t)\, e_t(\pi_t) \right] \end{aligned}$$

if we assume $\gamma_{-1}(h_{-1}, \pi_{-1}) = 0$. The first equality makes use of the fact that $p(h_{t+1}) = p(h_t)\, f(\pi_t|a_t(h_t, 1))\, \delta(\varsigma_{t+1}, \varsigma_t)$ and also combines $\beta^t \beta = \beta^{t+1}$. The second equality adds a term which is zero to make the discounted sum recursive in form; uses short-hand notation that $a_{\tau t} = a_t(h_t, \tau)$, $\gamma_{t-1} = \gamma_{t-1}(h_{t-1}, \pi_{t-1})$, $\gamma_t(\pi_t) = \gamma_t(h_t, \pi_t)$ and $e_t(\pi_t) = e_t(h_t, \pi_t)$;[25] uses the definition that $\rho_t = r_{t-1}(1, h_{t-1}, \pi_{t-1})$ along with the restriction that $1 - \rho_t = r_{t-1}(2, h_{t-1}, \pi_{t-1})$; and replaces probability sums with a conditional expectation.

The upshot is that we have the ability to create

$$\Psi_t = \left[-\gamma_{t-1}[\rho_t \mu(a_{1t}) + (1-\rho_t)\mu(a_{2t})] + \sum_{\pi_t \in \Pi} f(\pi_t|a_{1t})\, \gamma_t(\pi_t)\, e_t(\pi_t) \right] + \beta E_t \Psi_{t+1}$$

which indicates the importance of creating several "psuedo state variables" when we place this in recursive form, following the path of Marcet and Marimon (2010).

Combining the objective function and the Lagrangian component, the recursive

[25] We keep the argument π_t for γ_t and e_t to emphasize the fact that the inflation outcome π_t has not been realized when the policymaker is making decision about its inflation action at the beginning of period t.

form of the optimization problem is

$$(R1): W(\rho_t, \eta_t, \varsigma_t) = \min_{\gamma_t(\pi_t)} \max_{a_{1t}, e_t(\pi_t)} \left\{ w_t + \beta E_t W\left(\rho_{t+1}, \eta_{t+1}, \varsigma_{t+1}\right) \right\}$$

where the flow objective is

$$\begin{aligned} w_t &= \sum_{\pi_t \in \Pi} f(\pi_t | a_t(h_t, \tau = 1))\, u(\pi_t, e_t(\pi_t), \varsigma_t) \\ &\quad + \sum_{\pi_t \in \Pi} f(\pi_t | a_{1t})\, \gamma_t(\pi_t)\, e_t(\pi_t) - \eta_t \left[\rho_t \mu(a_{1t}) + (1-\rho_t)\mu(a_{2t})\right] \end{aligned}$$

and the state evolution equations are:

$$\begin{aligned} \rho_{t+1} &= \frac{\rho_t \theta(\pi_t, a_{1t})}{\rho_t \theta(\pi_t, a_{1t}) + (1-\rho_t)\theta(\pi_t, a_{2t})} \\ \eta_{t+1} &= \gamma_t \\ \text{prob}\left(\varsigma_{t+1} = s | \varsigma_t = \sigma\right) &= \delta(s, \sigma). \end{aligned}$$

A.3 The two-stage formation

The recursive form of the optimization problem (R1) has two stages. Define

$$W(\rho_t, \eta_t, \varsigma_t) = \max_{a_{1t}} \int \Omega(\rho_t, \eta_t, \varsigma_t, \pi_t) dF(\pi_t | a_{1t}) \tag{20}$$

where $F(\pi|a)$ is the distribution of inflation conditional on a particular policy action. And

$$\begin{aligned} \Omega(\rho_t, \eta_t, \varsigma_t, \pi_t) &= \min_{\gamma_t} \max_{e_t} \{ u(\pi_t, e_t, \varsigma_t) \\ &\quad + \gamma_t e_t - \eta_t \left[\rho_t \pi_t + (1-\rho_t) m\right] \\ &\quad + \beta E W(\rho_{t+1}, \eta_{t+1}, \varsigma_{t+1}) | \rho_t, \eta_t, \varsigma_t, \pi_t \} \end{aligned}$$

where m is short-hand for $E\pi_{t+1}|a_2$ in (18), which is predetermined to be $\mu + \phi\varsigma_t$ if $a_2 = \mu + \phi\varsigma_t$. If $a_2 = a_1 + \mu + \phi\varsigma_t$, m is still equal to $\mu + \phi\varsigma_t$ but the term in the squared bracket becomes:

$$l_t \pi_t + (1-\rho_t) m \text{ with } l_t = \rho_t + (1-\rho_t) = 1$$

After we substitute e_t by x_t using the equation (**??**), the $\Omega(\rho_t, \eta_t, \varsigma_t, \pi_t)$ becomes:

$$\begin{aligned}\Omega(\rho_t, \eta_t, \varsigma_t, \pi_t) &= \min_{\gamma_t} \max_{x_t} \{u(\pi_t, x_t) \\ &\quad + \gamma_t (\pi_t - \kappa x_t - \varsigma_t) - \eta_t [\rho_t \pi_t + (1 - \rho_t) m] \\ &\quad + \beta E W(\rho_{t+1}, \eta_{t+1}, \varsigma_{t+1}) | \rho_t, \eta_t, \varsigma_t, \pi_t\}\end{aligned} \tag{21}$$

which is the one in the main text.

B Explicit solution with fixed credibility

In this Appendix, we derive linear decision rules for the constant ρ models employed in sections 3 and 5 of the main text. We use a familiar procedure to do so: (i) finding the first order conditions; (ii) using the envelope theorem to eliminate the derivatives of the value function; and (iii) solving the resulting linear expectational difference system for the decision rules. The analysis does need, however, to take into account the two stage decision-making and the resulting rational expectations system also has a two stage structure.

Specifically, we seek the θ coefficients in the contingency plans

$$\begin{aligned}\gamma_t &= \underline{\gamma} + \theta_{\gamma\eta}(\eta_t - \underline{\eta}) + \theta_{\gamma\varsigma}\varsigma_t + \theta_{\gamma\varepsilon}\varepsilon_t \\ x_t &= \underline{x} + \theta_{x\eta}(\eta_t - \underline{\eta}) + \theta_{x\varsigma}\varsigma_t + \theta_{x\varepsilon}\varepsilon_t\end{aligned} \tag{22}$$

and in the optimal policy rule

$$a_t = \underline{a} + \theta_{a\eta}(\eta_t - \underline{\eta}) + \theta_{a\varsigma}\varsigma_t \tag{23}$$

To produce a solution that nests the automaton and tag-along cases, we specify expected inflation as

$$\begin{aligned}E_t^+ \pi_{t+1} &= \rho[E_t^+ a_{t+1}] + (1 - \rho)[\omega E_t^+ a_{t+1} + \mu + \phi E_t^+ \varsigma_{t+1}] \\ &= l[E_t^+ a_{t+1}] + (1 - \rho)\left(\mu + \phi E_t^+ \varsigma_{t+1}\right) \\ \text{with } l &= \rho + (1 - \rho)\,\omega\end{aligned}$$

The fixed-action alternative of section 3 corresponds to the case with $\omega = 0$ and the tag-along alternative of section 5 arises when $\omega = 1$.

The notation E_t^+ indicates that the expectation is conditional on π_t, while the notation E_t employed below corresponds to information at the start of a period (η_t, ς_t).

To ease derivations, we break the constraint

$$\pi_t - \kappa x_t - \beta E^+ \pi_{t+1} - \varsigma_t \geq 0$$

into two parts,

$$\begin{aligned} e_t - \beta E^+ \pi_{t+1} &\geq 0 \\ \pi_t - \kappa x_t - \varsigma_t &= e_t \end{aligned}$$

and use the latter to eliminate x and to write the momentary utility as

$$u\left(\pi, e, \varsigma\right) = -\frac{1}{2}\left[\pi^2 + q\left(\left(\pi - e - \varsigma\right) - \kappa x^*\right)^2\right]$$

We therefore also solve for θ coefficients for the variable e as part of the work.

B.1 The contingency plans and the intermediate value function

The monetary authority's contingency plans derive from maximization conditional on π_t and imply an intermediate value function Ω conditional on $\rho, \eta, \pi, \varsigma$. That is, the contingency plans for γ and e derive from the problem

$$\begin{aligned} P1 \quad : \quad & \Omega(\rho, \eta, \varsigma, \pi) = \min_{\gamma} \max_{e} \{u(\pi, e, \varsigma) \\ & + [-\eta(l\pi + (1-\rho)\left(\mu + \phi\varsigma\right)) + \gamma e] + \beta \widetilde{W}(\rho, \eta', \varsigma)\} \end{aligned}$$

where the maximization is subject to

$$\eta' = \gamma$$

and

$$\widetilde{W}(\rho, \eta', \varsigma) = EW(\rho, \eta', \varsigma') | \left(\eta', \varsigma\right)$$

The FOCs for this maximization are:

$$\begin{aligned} e &: q(\pi_t - e_t - \varsigma_t - \kappa x^*) + \gamma_t = 0 && (24) \\ \gamma &: e_t + \beta \frac{\partial}{\partial \eta'} \widetilde{W}(\rho, \eta' = \gamma_t, \varsigma) = 0 && (25) \end{aligned}$$

Applying the envelope theorem to the intermediate value function, we have that

$$\begin{aligned} \Omega_\pi(\rho, \eta, \pi, \varsigma) &= u_\pi(\pi, e, \varsigma) - \eta l = -\left[\pi + q\left((\pi - e - \varsigma) - \kappa x^*\right)\right] - \eta l && (26) \\ \Omega_\eta(\rho, \eta, \pi, \varsigma) &= -(l\pi + (1-\rho)(\mu + \phi\varsigma)) && (27) \end{aligned}$$

B.2 The policy action and the value function

The policy action and value function derive from problem

$$\begin{aligned} P2 &: W(\rho, \eta, \varsigma) \\ &= \max_a \int \Omega(\rho, \eta, \pi(a, \varepsilon), \varsigma) f(\varepsilon) d\varepsilon \end{aligned}$$

for which the FOC is

$$0 = \int \Omega_\pi(\rho, \eta, \pi(a, \varepsilon), \varsigma) f(\varepsilon) d\varepsilon$$

In these expressions and below, $\pi(a, \varepsilon) = a + \varepsilon$.

Using the ET result for the derivative of the intermediate value function, this simplifies to

$$0 = -\left[a + q\left(a - \int e\left(\rho, \eta, \pi(a, \varepsilon)\right) f(\varepsilon) d\varepsilon - \varsigma - \kappa x^*\right)\right] - \eta l \qquad (28)$$

.

The envelope theorem's application to the value function leads to

$$\begin{aligned} W_\eta(\rho, \eta, \varsigma) &= \int \Omega_\eta(\rho, \eta, \pi(a, \varepsilon), \varsigma) f(\varepsilon) d\varepsilon \\ &= -(la + (1-\rho)(\mu + \phi\varsigma)) \end{aligned}$$

where the second line involves use of the previous envelope theorem result for Ω_η.

B.3 Implications of the FOCs

To derive implications of the FOCs for the linear policy rules, we now introduce time subscripts into the equations above. With the envelope thereom results, the FOCs (24),(25) for e and γ can be written as

$$\begin{aligned} e &: q(\pi_t - e_t - \varsigma_t - \kappa x^*) + \gamma_t = 0 \\ \gamma &: e_t - \beta(lE_t^+ a_{t+1} + (1-\rho)\left(\mu + \phi E_t^+ \varsigma_{t+1}\right)) = 0 \end{aligned}$$

To faciliate the derivations, we start by solving these equations for e_t and γ_t.

$$\gamma_t = q\kappa x^* - q\pi_t + q\beta(lE_t^+ a_{t+1} + (1-\rho)\left(\mu + \phi E_t^+ \varsigma_{t+1}\right)) + q\varsigma_t \tag{29}$$

$$e_t = \beta(lE_t^+ a_{t+1} + (1-\rho)\left(\mu + \phi E_t^+ \varsigma_{t+1}\right)) \tag{30}$$

Turning to the FOC for a, we use the substitution $\eta_t = \gamma_{t-1}$ to write:

$$0 = -[a_t + q((a_t - \int e(\rho, \gamma_{t-1}, \pi(a_t, \varepsilon_t) f(\varepsilon) d\varepsilon) - \varsigma_t - \kappa x^*)] - \gamma_{t-1} l$$

and we then use the law of iterated expectations to write

$$\int e(\rho, \gamma_{t-1}, \pi_t = a_t + \varepsilon) f(\varepsilon) d\varepsilon = \beta(lE_t a_{t+1} + (1-\rho)(\mu + \phi E_t \varsigma_{t+1}))$$

Then the FOC for a becomes

$$(1+q)\, a_t = q\kappa x^* + q\beta(lE_t a_{t+1} + (1-\rho)(\mu + \phi E_t \varsigma_{t+1})) - \gamma_{t-1} l + q\varsigma_t \tag{31}$$

and a conditional expectation of (29) implies

$$E_t \gamma_t = q\kappa x^* - q a_t + q\beta(lE_t a_{t+1} + (1-\rho)(\mu + \phi E_t \varsigma_{t+1})) + q\varsigma_t \tag{32}$$

The four conditions (29),(30),(31), and (32) will next be use to determine the θ coefficients.

B.4 Solving for policy action and contingency plans

We first solve for the policy action and we then solve for the contingency plans. This reverses the natural ordering suggested by the dynamic programming analysis, but is simpler from the standpoint of linear rational expectations modeling.

B.4.1 The policy action

A solution to (31) and (32) is required to be of the form

$$\begin{aligned} E_t\gamma_t &= \underline{\gamma} + \theta_{\gamma\eta}(\gamma_{t-1} - \underline{\gamma}) + \theta_{\gamma\varsigma}\varsigma_t \\ a_t &= \underline{a} + \theta_{a\eta}(\gamma_{t-1} - \underline{\gamma}) + \theta_{a\varsigma}\varsigma_t \end{aligned}$$

The **steady-state coefficients** are readily shown to be

$$\underline{a} = (1-l)\,\underline{\gamma} = \frac{q\kappa x^* + q\beta(1-\rho)\mu}{(1-l)^{-1} + q\,(1-\beta l)}$$

The **feedback coefficients** $(\theta_{a\eta}, \theta_{\gamma\eta})$ are the solution to the equation system[26]:

$$\theta_{a\eta} = \theta_{\gamma\eta} - l \tag{33}$$

$$\theta_{\gamma\eta} = -q\theta_{a\eta} + \beta q l \theta_{a\eta}\theta_{\gamma\eta} \tag{34}$$

Substituting (33) into (34) leads to

$$\theta_{\gamma\eta} = -q(\theta_{\gamma\eta} - l) + \beta q l(\theta_{\gamma\eta} - l)\theta_{\gamma\eta}$$

and we thus $\theta_{\gamma\eta}$ to be the stable root of

$$\beta q l\ z^2 - [\beta q l^2 + (1+q)]\ z + ql = 0$$

[26] It is conventional to discuss this solution in terms of an eigenvalue. Substituting (33) into (34) leads to

$$\theta_{\gamma\gamma} = -q(\theta_{\gamma\gamma} - l) + \beta q l(\theta_{\gamma\gamma} - l)\theta_{\gamma\gamma}$$

and we thus take $\theta_{\gamma\gamma}$ to be the stable root of

$$\beta q l\ z^2 - [\beta q l^2 + (1+q)]\ z + ql = 0$$

Then, (33) delivers $\theta_{a\gamma} = \theta_{\gamma\gamma} - l$.

Then, (33) indicates $\theta_{a\eta} = \theta_{\gamma\eta} - l$.

To derive the **response coefficients**, we suppose that

$$E_t \varsigma_{t+1} = \delta \varsigma_t$$

so that

$$\begin{aligned} E_t a_{t+1} &= \underline{a} + \theta_{a\eta}(E_t \gamma_t - \underline{\gamma}) + \theta_{a\varsigma} E_t \varsigma_{t+1} \\ &= \underline{a} + \theta_{a\eta}(E_t \gamma_t - \underline{\gamma}) + \theta_{a\varsigma} \delta \varsigma_t \end{aligned}$$

Then, (31) and (32) imply

$$\begin{aligned} \theta_{\gamma\varsigma} &= \theta_{a\varsigma} \\ \theta_{\gamma\varsigma} &= -q\theta_{a\varsigma} + \beta q l \theta_{a\eta} \theta_{\gamma\varsigma} + \beta q l \theta_{a\varsigma} \delta + \beta q \left(1 - \rho\right) \phi \delta + q \end{aligned}$$

so that

$$\theta_{\gamma\varsigma} = \theta_{a\varsigma} = \frac{\beta q \left(1 - \rho\right) \phi \delta + q}{1 + q - q\beta l \left(\theta_{a\eta} + \delta\right)}$$

We now know the full set of policy action coefficients $\underline{a}, \theta_{a\eta}, \theta_{a\varsigma}$ and we know $\underline{\gamma}, \theta_{\gamma\eta}, \theta_{\gamma\varsigma}$.

B.4.2 The contingency plans

Note further that

$$E_t^+ a_{t+1} = \underline{a} + \theta_{a\eta}(\gamma_t - \underline{\gamma}) + \theta_{a\varsigma} \delta \varsigma_t$$

so that subtracting (32) from (29) lets us directly determine the response of the multiplier to an implementation error

$$\begin{aligned} \gamma_t - E_t \gamma_t &= -q\varepsilon_t + \beta q l (E_t^+ a_{t+1} - E_t a_{t+1}) \\ &= \frac{-q}{1 - \beta q l \theta_{a\eta}} \varepsilon_t = \theta_{\gamma\varepsilon} \varepsilon_t \end{aligned} \tag{35}$$

The behavior of expected information is then easy to determine in parallel form using (30)

$$\begin{aligned}
e_t &= \beta(lE_t^+ a_{t+1} + (1-\rho)\left(\mu + \phi E_t^+ \varsigma_{t+1}\right)) \\
&= \beta l[\underline{a} + \theta_{a\eta}(\gamma_t - \underline{\gamma}) + \theta_{a\varsigma}\delta\varsigma_t] + \beta(1-\rho)\left(\mu + \phi\delta\varsigma_t\right) \\
&= \beta l\underline{a} + \beta(1-\rho)\mu \\
&\quad + \beta l\theta_{a\eta}\theta_{\gamma\eta}(\eta_t - \underline{\eta}) \\
&\quad + \beta\left[l\left(\theta_{a\eta}\theta_{\gamma\varsigma} + \theta_{a\varsigma}\delta\right) + (1-\rho)\phi\delta\right]\varsigma_t \\
&\quad + \beta l\theta_{a\eta}\theta_{\gamma\varepsilon}\varepsilon_t
\end{aligned}$$

thus determining the coefficients $\underline{e}$, $\theta_{e\eta}$, $\theta_{e\varsigma}$ and $\theta_{e\varepsilon}$.

Finally, given that $\pi_t - \kappa x_t - \varsigma_t = e_t$, the coefficients $\underline{x}$, $\theta_{x\eta}$, $\theta_{x\varsigma}$ and $\theta_{x\varepsilon}$ must satisfy

$$\begin{aligned}
\underline{x} &= \frac{1}{\kappa}[\underline{a} - \underline{e}] \\
\theta_{x\eta} &= \frac{1}{\kappa}[\theta_{a\eta} - \theta_{e\eta}] \\
\theta_{x\varsigma} &= \frac{1}{\kappa}[\theta_{a\varsigma} - \theta_{e\varsigma} - 1] \\
\theta_{x\varepsilon} &= \frac{1}{\kappa}[-\theta_{e\varepsilon} + 1]
\end{aligned}$$

B.5 The inflation action in terms of observables

Notice that (23) implies:

$$(\eta_{t-1} - \underline{\eta}) = \frac{1}{\theta_{a\eta}}(a_{t-1} - \underline{a}) - \frac{\theta_{a\varsigma}}{\theta_{a\eta}}\varsigma_{t-1}$$

Insert the result into (22), use the fact that $\eta_t = \gamma_{t-1}$ and then solve it into the policy rule

$$\begin{aligned}
\eta_t - \underline{\eta} &= \theta_{\gamma\eta}(\eta_{t-1} - \underline{\eta}) + \theta_{\gamma\varsigma}\varsigma_{t-1} + \theta_{\gamma\varepsilon}\varepsilon_{t-1} \\
&= \frac{\theta_{\gamma\eta}}{\theta_{a\eta}}(a_{t-1} - \underline{a}) + [\theta_{\gamma\varsigma} - \frac{\theta_{\gamma\eta}\theta_{a\varsigma}}{\theta_{a\eta}}]\varsigma_{t-1} + \theta_{\gamma\varepsilon}\varepsilon_{t-1}
\end{aligned}$$

and then substitute it into (23):

$$\begin{aligned} a_t &= \underline{a} + \theta_{\gamma\eta}(a_{t-1} - \underline{a}) \\ &\quad + [\theta_{a\eta}\theta_{\gamma\varsigma} - \theta_{\gamma\eta}\theta_{a\varsigma}]\varsigma_{t-1} + \theta_{a\eta}\theta_{\gamma\varepsilon}\varepsilon_{t-1} \\ &\quad + \theta_{a\varsigma}\varsigma_t \end{aligned}$$

B.6 Discretionary rule

Using the expressions above, the discretionary rule $\mu + \phi\varsigma$ should be the solution to the partial credibility case with $\rho = 0$ and $\omega = 0$. That is

$$\mu + \phi\varsigma = \underline{a} + \theta_{a\varsigma}\varsigma$$

where $\underline{a}$ and $\theta_{a\varsigma}$ are evaluated at $\rho = 0$ and $l = 0$. This determines the coefficients:

$$\begin{aligned} \mu &= \frac{q\kappa x^*}{1 + q - \beta q} = \frac{\kappa h x^*}{\kappa^2 + (1 - \beta)\, h} \\ \phi &= \frac{q}{1 + q - \beta q\delta} = \frac{h}{\kappa^2 + (1 - \beta\delta)\, h} \end{aligned}$$

Those coefficients are consistent with the solution derived directly from the maximiztion with a discretionary monetary authority.

C Calibration

This Appendix explains our calibration strategy.

C.1 A mapping from the textbook loss function to the objective function in the paper

In Chapter 5 Appendix 1 of Gali's textbook, a second order approximation to the consumer's welfare losses yields a quadratic loss function for the central bank:

$$\sum_{t=0}^{\infty} \beta^t [\frac{1}{2}\frac{\varepsilon}{\lambda}\pi_t^2 + \frac{1}{2}[\sigma + \frac{\gamma + \alpha}{1 - \alpha}]\widehat{x}_t^2 - \Phi\widehat{x}_t]$$

where σ, γ are coefficients in the utility function $\frac{1}{1-\sigma}c^{1-\sigma} - \frac{1}{1+\gamma}n^{1+\gamma}$; α is the share of fixed factor in production function ($y = an^{1-\alpha}$); ε is the demand elasticity; $\Phi = 1 - \frac{1}{M}$ and $M = \frac{\varepsilon}{\varepsilon-1}$ so that $\Phi = \frac{1}{\varepsilon}$; and λ is the slope coefficient obtained from estimating the Phillips Curve using marginal cost proxy.

$$\pi_t = \beta E_t \pi_{t+1} + \lambda(\psi_t - \psi)$$

Some algebra enables us to rewrite the quadratic loss function as follows:

$$\sum_{t=0}^{\infty} \beta^t [\frac{1}{2}\frac{\varepsilon}{\lambda}\pi_t^2 + \frac{1}{2}A(\widehat{x}_t - \frac{\Phi}{A})^2 - \frac{1}{2}A(\frac{\Phi}{A})^2$$

with $A = [\sigma + \frac{\gamma+\alpha}{1-\alpha}]$. $\widehat{x}_t$ in the loss function is defined as:

$$\widehat{x}_t \equiv \widehat{y}_t - \widehat{y}_t^e = (y_t - y) - (y_t^e - y^e)$$

where hat means the deviation from the steady state and y_t^e denotes the efficient level of output.

In our paper, the objective function of the central bank is assumed to be

$$\sum_{t=0}^{\infty} -\frac{1}{2}[\pi_t^2 + h(x_t - x^*)^2]$$

where $x_t \equiv y_t - y_t^n$ is the output gap that also enters the Phillips Curve. It can be shown that $\widehat{x}_t$ is equal to x_t and therefore, there is a direct mapping from the textbook quadratic loss function to the objective function that we assume in our paper::

$$\begin{aligned} h &= A\frac{\lambda}{\varepsilon} \\ x^* &= \frac{\Phi}{A} \end{aligned}$$

In addition, Chapter 3 in Gali's textbook gives the connection between the real marginal cost and the output gap: $(\psi_t - \psi) = Ax_t$, so that we have a mapping from the estimated PC to our functional form PC: $\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + \varsigma_t$ where

$$\kappa = \lambda A$$

C.2 The baseline calibration

In the literature, there are some popular estimates for λ. We take the one from Gali-Gertler (1999) where $\lambda = .04$, if inflation is measured as the quarterly change in the price level and β is constrained close to one. There is much less agreement about the elasticity of marginal cost with respect to the output, i.e., $A = [\sigma + \frac{\gamma+\alpha}{1-\alpha}]$. We set $A = 1$ following Van Zandweghe and Wolman (2011), in which they assume $\sigma = 1$, $\gamma = 0$ and $\alpha = 0$. A reasonable value for the demand elasticity ε is in the range of 4 (implying a gross markup of M 1.33) to 10 (implying $M = 1.1$) and we set it to be 10.

These values of deep parameters determines the coefficients in the objective function and the Phillips Curve:

$$h = .004;\ x^* = .2;\ \kappa = .04.$$

We take the relevant period to correspond to a quarter. We assume $\beta = 0.995$, which implies a steady state interest rate about two percent. The implied inflation bias and stablization bias are:

$$\begin{aligned}
\mu &= \frac{\frac{\kappa}{\varepsilon}}{(1-\beta+\frac{\kappa^2\varepsilon}{\lambda})} = \frac{\kappa h x^*}{\kappa^2+(1-\beta)\,h} = 1\% \\
\phi &= \frac{1}{\varepsilon\kappa+1-\beta\delta} = \frac{h}{\kappa^2+(1-\beta\delta)\,h} = 1.98
\end{aligned}$$

The inflation shock ς_t in the Phillips Curve is assumed to have presistence .9 and a standard deviation 2%. Finally the standard deviation of the implementation error is assumed to be 1%.

D Computation

When there is learning via Bayes's law, the model is no longer linear quadratic. For this reason, we use polynomial approximations to the value functions and fine grids for the decision variables. The next version of this paper will describe the computational method in detail.