Optimal policy design with a skeptical forward-looking private sector*

Robert G. King† and Yang K. Lu‡

Feb 14, 2013

Abstract

How should a committed policymaker optimally design his plans when private agents are skeptical that these plans will be carried out? We pose this policy question in a reputation game between policymakers of various types and private agents, and in a setting with private agents being forward-looking, as is standard in macroeconomics. We provide a recursive formulation of the optimal policy problem, whose constraints involve rational Bayesian learning by private agents and incentive constraints associated with regulating the behavior of an alternative opportunistic policymaker. This recursive method makes it feasible to study the dynamics of reputation effects on optimal policy.

(This version is incomplete, providing an overview of the dynamic theory but no computational results.

Very preliminary. Please do not circulate without authors' permission.)

*The authors would like to thank Michael Golosov, Albert Marcet, Ramon Marimon, Nicola Pavoni, Chen Wang, seminar participants at Fudan University for helpful comments, and the Fondation Banque de France for financial support.

†Boston University

‡Hong Kong University of Science and Technology

"From this evidence of a favorable disposition, given by the former government, the institution of a new one, clothed with powers competent to calling forth the resources of the community, has excited correspondent expectations. A general belief, accordingly, prevails, that the credit of the United States will quickly be established on the firm foundation of an effectual provision for the existing debt. The influence, which this has had at home, is witnessed by the rapid increase, that has taken place in the market value of the public securities...." Alexander Hamilton, Report on the Public Credit, January 9 1790.

"we have not followed the theoretical prescription of promising to keep rates low enough for long enough to create a period of above-normal inflation. The arguments in favor of such a policy hinge on a clear understanding on the part of the public that the central bank will tolerate increased inflation only temporarily–say, for a few years once the economy has recovered–before returning to the original inflation target in the long term. In standard theoretical model environments, long-run inflation expectations are perfectly anchored. In reality, however, the anchoring of inflation expectations has been a hard-won achievement of monetary policy over the past few decades, and we should not take this stability for granted. Models are by their nature only a stylized representation of reality, and a policy of achieving "temporarily" higher inflation over the medium term would run the risk of altering inflation expectations beyond the horizon that is desirable. Were that to happen, the costs of bringing expectations back to their current anchored state might be quite high." Donald L. Kohn, "Monetary Policy Research and the Financial Crisis: Strengths and Shortcomings", October 9, 2009.

1 Introduction

How should a committed policymaker optimally design his plans when private agents are skeptical that these plans will be carried out? This question has confronted policymakers for centuries, as the quotations from Alexander Hamilton in 1790 and Donald Kohn in 2009 illustrate. The early United State decision to fully fund the debt from the prior confederation of revolutionary colonies was a difficult step that build substantial credibility for the young country. The Federal Reserve System, concerned about its anti-inflation credibility, chose to act more modestly in specifying inflation plans during the recent financial crisis than one have been called for by optimal commitment policies in the absence of private sector skepticism. Each of these quotation suggests that the interplay between policy actions and expectations is important in fiscal and monetary arenas, but that desirable actions must take into account effects on expectations and beliefs.

In this paper, we provide a basic forward-looking economy in which our introductory question about policy is well-posed and, which in broad form, the concerns of Hamilton and Kohn are highlighted. There is a precise answer for a policymaker who is trustworthy, in the sense that his future actions will be his current announcements of plans. Calculation of this optimal policy is made feasible by the derivation of a recursive optimal policy problem, whose constraints involve rational Bayesian learning by private agents and incentive constraints associated with regulating the behavior of an alternative opportunistic policymaker. The recursive optimum problem and the associated optimal policies each involve the management of expectations in a central manner.

Much recent macroeconomic research has been devoted to developing theoretical and empirical implications of optimal policy design in forward-looking environments with complete commitment or no commitment. By contrast, our research program seeks to develop theoretical and empirical implications of optimal policy in the middle ground where many policymakers actually find themselves, with an internal ability to commit to a degree (even fully as we assume) but also with private agents are learning if he is such a type or has less commitment ability (even none). In the larger research program, we aim to create both illustrative examples and operational small-scale macroeconomic models, with the latter requiring development of complementary computational methods. In the current paper, we develop a sharply focused example that highlights the essential ideas which arise in this program.

There are four key features of the example environment. First, economic outcomes in a given period (t) depend on expectations about economic outcomes in the next period (t+1), which in turn depend on policy actions at t+1: it is in this sense that our policy authority faces a forward-looking private sector. Second, the type of the policymaker is unknown to private agents, private agents plausibly may have skepticism about an announced policymaker course of action. In particular, in addition to the committed or strong policy maker, who is the focus of our policy design problem, there is a second type whose actions are chosen on a period-by-period basis. Third, policymaker actions are unobservable and these have a random impact on economic outcomes. Fourth, there is private sector learning about policymaker type from observations on economic outcomes.

Our work thus is related to five literatures. The first is research that employs policy design under commitment along Ramsey [1927] lines, which has been applied in a wide range of macroeconomic settings.[1] The second is macroeconomic analysis of environments with private sector learning about unobserved aspects of policy rules.[2] The third is research on repeated games with reputational effects and, more specifically, work on environments in which actions are unobserved, but can be partially inferred from observable outcomes, which is typically described as dynamic games with imperfect public monitoring.[3] The fourth is research on optimal contracts, particularly dynamic contracts in recursive form, and the branch of this literature that uses a first order condition approach.[4] The fifth is the study of signalling equilibria and the related macroeconomic research on transparency and communication.

Relative to these diverse literatures, our contribution is to provide an approach to policy design by the committed type given that private agents are forming expectations based on a common public history of outcomes and given that private agents correctly understand the actions that could be taken by the alternative, opportunistic type as well as the committed type. Notably, this approach involves the construction of an optimal contract, in which the opportunistic type's optimal behavior is summarized by a first order condition and a value function. By adopting this research strategy in this initial draft, we sidestep the issue of the nature of the signalling game between the types of the policy authority and the private sector, but we conjecture that the constrained Pareto

[1] Exemplified by analyses of optimal taxation over time (as in Chamley [19xx] and Lucas and Stokey's [19xx]).

[2] Discuss work originating in Taylor

[3] Use Mailath survey to set focus

[4] add references

optimality between signal senders will mean that the outcomes we study are those which uniquely satisfy the refinement criteria of Mailath et. al. [1993] We take this approach because of the general close link between signalling equilibria that satisfy the Pareto criterion and refinement based on applying Bayesian reasoning to the interpretation of out-of-equilibrium signals. In particular, we have found this to be the case in earlier work on dynamic repeated games with a similar structure (King, Lu and Pasten 2008).

One of our key motivations for developing this approach is to study the dynamics of reputation effects on optimal policy. From much research on the boundaries of game theory and macroeconomics, it is well understood that, in an environment of imperfect public monitoring, there are strong reputational forces that will affect the actions of an opportunistic policymaker if the private sector believes that there is a probability that there is a committed type which plays a fixed action. In the early periods of a finite horizon game, many models suggest that the opportunistic type will take actions close to the commitment action, so as to avoid private sector learning. For example, Backus and Driffill [1985a,b] and Barro [1986] provide monetary policy analyses that display this phenomenon, building on the insights of Kreps and Wilson [1982b] and Milgrom and Roberts [1982]. In the later stages of these finite horizon repeated games, reputation effects are less influential and learning phenomena more pronounced. Consequently, when there is policymaker replacement after a fixed number of periods, the early stages of "regimes" should be marked by substantial discipline on opportunistic types from reputation and limited learning, which would dissipate over time. More generally, Phelan's analysis of a repeated game with observable tax actions shows that similar phenomena arise when type changes (replacements) are unobserved. Starting from an initial condition in which the likelihood of a committed type is low, the opportunistic type will have a relatively strong incentive (in a probability sense) to mimic the strong type.

By contrast, in the settings that are of interest to us, the intuition is that the committed type will select its policies to induce learning in the early stages of a regime, particularly if it is more patient. In a dynamic repeated game with observable tax actions and fixed unobservable type, Lu [2013] shows that the committed type induces the opportunistic type to play mixed strategies which lead to rapid private sector learning about type in the early stages of the game. The intuition is that the committed type takes tough actions in the early stages of the game as an investment in a form of credibility capital, i.e., in a reputation for being the type of government that does not optimally engage in confiscatory taxation.

However, in settings with a forward-looking private sector as are standard in macroeconomics, optimal policies under commitment typically feature history dependence and, specifically, an initial period in which a policymaker optimally exploits private sector initial conditions. Such "start up" phenomena are pervasive: they include initial capital levies or high rates of taxation of income, initial intervals of high inflation to tax real balances or exploit the Phillips curve, absence of enforcement of intellectual property rights, and so on. However, when the ability to commit to desirable long-term policy is uncertain, then these intrinsic incentives for a committed policymaker to have a history-dependent plan that treats early periods differently must be considered jointly with reputation effects, i.e., with the incentive to accumulate credibility capital.

The organization of the paper is as follows. Section 2 lays out the commonly employed macroeconomic model, which has the attractive property that it includes forward-looking behavior without any endogenous state variables, which allows our analysis to focus on other state-like variables introduced by commitment and learning. Section 3 defines and characterizes the equilibrium of this model. Section 4 proposes a recursive formulation of the policy problem that can be employed to compute the equilibrium of the dynamic game. (Appendix A provides details of our derivations). Section 5 revisits issues for the policy announcements and shows that the solution to the recursive policy problem described in Section 4 is the unique equilibrium announcement surviving the refinement in the spirit of Mailath et. al (1993). Section 6 concludes.

2 Setup

There has been much recent interest in designing policy rules under commitment that maximize an expected present discounted value objective, subject to a series of one or more forward-looking constraints. In this section, we describe a similar economy in which there is a committed policymaker and private agents are forward-looking. However, in our economy, private agents have some skepticism about whether an announced plan will in fact be executed.

2.1 Private sector

There is a *private sector*, which is composed by all agents in the economy but the policymaker. A key characteristic of agents in the private sector is that none of them has strategic power (they are atomistic). A reduced-form of the aggregate outcome of the interaction among agents in the private sector is represented by the function

$$x_t = f(\pi_t, e_t, \varsigma_t) \tag{1}$$

where the aggregate endogenous variable x_t depends on the policy outcome π_t,[5] the forward-looking variable e_t that is defined as[6]

$$e_t = \beta \mathbb{E}_t \pi_{t+1}, \tag{2}$$

the private sector's discounted expectations of future policy outcome conditioned on its information in period t (the private sector's conditional expectation is denoted by $\mathbb{E}_t$), and the exogenous variable ς_t which is assumed to follow a Markov process

$$prob(\varsigma_t = s | \varsigma_{t-1} = \sigma) = \delta(s, \sigma). \tag{3}$$

2.2 Policymakers: types and actions

There is also a *policymaker* who lives forever and may be one of two types: *strong* (type $\tau = 1$) with probability g or *opportunistic or weak* (type $\tau = 2$) with probability $1 - g$. Both types of policymakers must choose policy actions $\{a_t\}_{t=0}^{\infty}$ to be implemented at each period t. Policy actions affect policy outcomes $\{\pi_t\}_{t=0}^{\infty}$ contemporaneously according to the distribution $\Theta(\pi_t \mid a_t)$. Policymakers' types and policy actions are not directly observable by the private sector; only policy outcomes are publicly observable once realized. Either type has to make announcements $\{\alpha_t\}_{t=0}^{\infty}$ about his policy actions at $t = 0$. The main differences between two types of policymakers are the timing in which they decide policy actions and their freedom to set announcements at $t = 0$.[7]

[5] We make a distinction between policy actions and policy outcomes which specified below.

[6] The scalor variables may be extend to be vectors and the expectation may be more general than the linear function here.

[7] A sentence may be modified or removed later.

Strong policymaker. The strong type ($\tau = 1$) is interpreted to be endowed with a commitment technology. He chooses the complete sequence of policy actions $\{a_t(\tau = 1)\}_0^\infty \in A$ at $t = 0$ which takes the form of a policy plan contingent on public history. In addition, this type is restricted to be *truthful* or *trustworthy* in the sense that announcements must coincide with the policy plan, $\widetilde{a}_t(\tau = 1) = a_t(\tau = 1)\ \forall t$.

The momentary objective of the strong policymaker is $v(\pi_t, x_t, \varsigma_t)$, which depends on the policy outcome π_t, the aggregate endogenous variable x_t delivered by the interaction of private agents, and the exogenous variable ς_t. Using equation (1) to replace x_t in the function $v(\cdot)$, the objective of the strong policymaker becomes:

$$E(U_t|H_t) = E\{\sum_{s=t}^{\infty} \beta^{s-t} u(\pi_s, e_s, \varsigma_s)|H_t\} \tag{4}$$

where $E(\cdot|H_t)$ is the strong policymaker's expectation based on his information set at period t, H_t, e_s is defined as in (2), and the discount factor β is assumed to coincide with the discount factor of agents in the private sector.

Weak policymaker. The opportunistic or weak type ($\tau = 2$) is interpreted to lack commitment. He chooses its policy actions $\{a_t(\tau = 2)\}_0^\infty \in A$ sequentially every period t, so we sometimes speak of a series of weak policymakers taking actions on a period-by-period basis. In addition, the weak type, if in place, is free to set its announcements at $t = 0$. Each such weak policymaker knows that the private sector cannot observe his actions but their expectations are sensitive to economic outcomes. Without loss of generality and to permit comparisons to other research, we allow the weak policymaker to differ in his discount factor (writing it as b, relative to the common discount factor of β for the strong type and the private sector) and in momentary objective $\widetilde{u}(\pi_t, x_t, \varsigma_t)$. Hence, his objective is[8]

$$E(\widetilde{U}_t|\widetilde{H}_t) = E\{\sum_{s=t}^{\infty} b^{s-t} \widetilde{u}(\pi_s, e_s, \varsigma_s)|\widetilde{H}_t\} \tag{5}$$

where $E(\cdot|\widetilde{H}_t)$ is the weak policymaker's expectation based on his information set at period t, $\widetilde{H}_t$, and e_s has the same definition as in (2).

[8] [Yang: Shall we use another expectation notation since the information set of weak policymaker is difference from the strong policymaker]

2.3 Timing of actions

At the beginning of period $t = 0$, the type τ of the policymaker is drawn and then the policymaker in place makes its announcements $\{\alpha_t(\tau)\}_{t=0}^{\infty}$.

After the announcements, every period $t \geqslant 0$ is broken into the following stages:[9]

(0) At the start of the period, the exogenous variables ς_t are drawn and observed by both the private sector and the policymaker. Hence, all agents observe public history $h_t = \{h_{t-1}, \pi_{t-1}, \varsigma_t\}$, which is composed by all past realizations of policy outcomes and exogenous shocks, together with the current shock ς_t. History also includes the outcome of the interaction of the private sector x_t, but we abstract from it since x_t is summarized by the policy outcome π_t, the exogenous variable ς_t and expectations e_t, which also depends on history. We thus use h_t to denote the public's information set (history) at the *start of the period*.

(1-a) If he is present, the strong policymaker follows the announced policy action represented by a policy plan contingent on the public history, $a_t(h_t, \tau = 1) \in A$.

(1-b) If he is present, the weak policymaker takes an action at each period t based on the public history h_t, $a_t(h_t, \tau = 2) \in A$.

In either case, the policymaker's action is unobservable to private agents. However, the private information of his past policy actions does not matter for the policymaker's decision-making. We will come back to this restriction later in discussing the equilibrium of the game. In addition, two types of policymakers share the same set of feasible policy actions A.

(2) The policy action a_t results in a random policy outcome π_t, according to

$$prob(\pi_t = z | a_t = a) = \theta(z, a) \tag{6}$$

with policy action $a_t \in A$ and policy outcome $\pi_t \in \Pi$.

[9] We need to decide whether we want the policy game to start in $t = 0$, so the stages below are also valid for $t = 0$, or start in $t = 1$ letting $t = 0$ only for announcements—in which case sums in the objective functions above should start with $j = 1$.

Does the contingent plans have to be conditional on knowing s(0)? – Risk aversion

If so, then the announcement has to be in the same period of the action

(3) Based on the history h_t and the observed policy outcome π_t, private agents update their belief that a policymaker of type $\tau = 1$ is present, which takes the general form for each type

$$prob(\tau = 1|h_t, \pi_t) \equiv r_t(\tau = 1, h_t, \pi_t). \tag{7}$$

The definitional equality in (7) arises because it is convenient for us to have a notation for how private agents' belief about type responds to new information, π_t, for which we will use $r_t(\tau, h_t, \pi_t)$ as above. In contrast to the policy outcome π_t, the exogenous shock ς_{t+1} does not contain any information about the type of the present policymaker. Thus, $prob(\tau = 1|h_t, \pi_t) = prob(\tau = 1|h_{t+1})$.

Since there are only two types, there is only one independent probability. We will focus on private agents' belief that a strong type is present as a key state variable in the analysis below, adopting the notation that $\rho_t \equiv prob(\tau = 1|h_t)$. It is a predetermined variable with respect to π_t. Correspondingly, private agents' belief that a weak type is present is $1 - \rho_t = prob\,(\tau = 2|h_t)$.

Combining these notations, the key state variable will evolve according to

$$\rho_{t+1} = r_t(1, h_t, \pi_t)$$

(4) Conditional on the updated belief about type, the private agents then form expectations of next-period policy outcome $E\pi_{t+1}|\left\{h_t, \pi_t\right\}$.

$$E\pi_{t+1}|\left\{h_t, \pi_t\right\} = \sum_{\tau} r_t\left(\tau, h_t, \pi_t\right) \sum_{\varsigma_{t+1}\in S} \delta(\varsigma_{t+1}, \varsigma_t)\mu(\bar{a}_{t+1}(h_{t+1}, \tau)) \tag{8}$$

where $\bar{a}_{t+1}(h_{t+1}, \tau)$ is private agents' conjectured policy action of type τ policymaker next period. We form discounted expectations $e_t\left(h_t, \pi_t\right) = \beta E\pi_{t+1}|\left\{h_t, \pi_t\right\}$ as a model element.

In (8) and elsewhere in the paper, we use

$$\mu(a) = \sum_{\pi\in\Pi} \pi\theta(\pi, a)$$

to denote the mean of policy outcome which will occur if the policy action is a.

Notice that when private agents form expectations, they take the policymaker actions in periods t and t+1 as given because each private agent is atomistic.

2.4 Applications

We now precise some environments that may be nested in the generic model describe above.

A New Keynesian Example. One example fits into this framework is the well-known New Keynesian model as variously elaborated in Clarida, Gali and Gertler [xxxx], Gali [xxxx], Woodford [xxxx] and Walsh[xxxx]. In this example, the policymaker is the monetary authority and the policy outcome (π) is the inflation. The momentary objective for the monetary authority is assumed to depend on both the inflation (π) and the output gap (x). A series of forward-looking constraints are present because the output gap is determined period-by-period through the New Keynesian Phillips curve

$$\pi_t = \alpha x_t + \beta \mathbb{E}_t \pi_{t+1} + \varsigma_t \tag{9}$$

which includes an important forward-looking component – the private sector's expected future inflation ($e_t = \beta \mathbb{E}_t \pi_{t+1}$), as well as the inflation (π_t) and a fundamental "price shock" (ς_t) governed by a Markov process. In the near-term drafts of this paper, we stick to this example in our computational exercise. However, we view the general ideas as applicable to a wide range of policy problems as discussed in the introduction above and we now provide a further discussion of several.

Revenue from money creation. Consideration of the classic issue of the revenue from money creation helps indicate that our general ideas are applicable to a wide range of policy problems, as well as highlighting some simplifying aspects of our example. Suppose that the policymaker is concerned with maximizing the expected present discounted value of the revenue from money creation,

$$\frac{M_t - M_{t-1}}{P_t} = m_t - \frac{1}{\pi_t} m_{t-1}$$

where M_t is the nominal stock of money, P_t is the price level, m_t is the real stock of money, and $\pi_t = P_t / P_{t-1}$ is the inflation rate.

Suppose further that the private sector's real demand for money is implicit in a first-order condition

$$f(m_t, \varsigma_t) = \mathbb{E}_t(\frac{1}{\pi_{t+1}})$$

as could arise in variety of monetary models. Note that the objective now contains a lag of real balances and that both sides of the constraint involve nonlinear functions. Neither of these is an important problem for the approach. The inclusion of a real state variable into the recursive objective is direct along dynamic programming lines. The analysis of Marcet and Marimon [2010], on which we build, allows for forward-looking constraints of the form $f_1(m_t, e_t, \varsigma_t) = \beta E_t f_2(m_{t+1}, e_{t+1,}\varsigma_{t+1})$ in a full information setting, so that it is not difficult to extend the example so as to accommodate the nonlinearity.

Substantively, a well-known feature of maximization of the revenue from money creation is that it is optimal for the initial period inflation rate to used to effectively confiscate the pre-existing real balances when there is full commitment, while setting a more moderate rate of inflation in subsequent periods. Yet, an optimizing government in our framework would presumably set a much lower rate of inflation, so as to distinguish itself from the alternative opportunistic government.

The Public Debt. The theory of fiscal policy with public debt contains a similar puzzle: why would a new, optimizing, committed government not renege on the public debt of a prior regime? A version of our analysis could provide insight into the dynamics of debt with imperfect credibility. The action to honor the preexisting debt, like that taken by the first U.S. treasury secretary Alexander Hamilton, could be part of an optimal policy to build fiscal credibility at an optimal rate.

3 The Equilibrium

In the previous section, we lay out the structure of a dynamic game between two types of policymakers and private agents. This section analyzes the strategic interaction among players and defines equilibrium in this dynamic game. The analysis highlights the interplay between the strong policymaker's actions and the learning behavior of private agents, in light of the potential actions that would have been taken if the policymaker in present were weak.

3.1 Equilibrium in policy announcements

The part of the game in which policymakers choose announcements at $t = 0$ may be studied separately with respect to the game in which policymakers choose policy actions. At the present stage of presentation, we simply assume that there is a unique equilibrium in announcements with a "leader-follower" structure. The strong type is the leader by announcing whatever is its optimal policy plan, and the weak type is the follower by announcing the same plan. Thus,

$$\alpha_t(h_t, \tau = 2) = a_t(h_t, \tau = 1)$$

and private agents learn nothing about the government type from observing the announcement:

$$\rho_0 = g.$$

In Section 5 of this paper, we will come back to this assumption and show that it is indeed the unique equilibrium in announcements that survives the refinement of Mailath et al [1993].

3.2 Equilibrium in policy actions

3.2.1 Equilibrium concept: Perfect Bayesian Public Equilibria

As some readers may have already noticed, when we specify the actions of two types of policymakers, those actions are solely conditional on the publicly observable history h_t. That is, $a_t(h_t, \tau = 1, 2)$ are "public strategies," where policymakers ignore their private information about their own past actions in choosing their actions. In turn, we restrict our attention to equilibria in "public strategies", namely "public equilibria".

We further require the equilibrium in this imperfect information game to be perfect Bayesian.

3.2.2 Consistent belief and rational expectations of the private agents

It is important to notice the "imperfect public monitoring" feature in our economy. That is, the policy actions are not directly observable to private agents. Instead, they determine the likelihood of various policy outcomes that are observable so that private agents may draw statistical inference on how likely a particular outcome comes from the policy action of the type $\tau = 1$ policymaker.

In order to make such an inference, private agents need to know the policy actions of both types of policymakers. The strong policymaker's actions are known since they will be determined fully by the announced strategies at time $0 : \{a_t(h_t, \tau = 1)\}_{t=0}^{\infty} = \{\alpha_t(h_t, \tau = 1)\}_{t=0}^{\infty}$. As for the policy action of the weak policymaker at period t, private agents will have to guess it.

Therefore, the learning functions $\{r_t(\tau, h_t, \pi_t)\}_{t=0}^{\infty}$ and the expectation functions $\{e_t(h_t, \pi_t)\}_{t=0}^{\infty}$ are based on the announced strategy profile of the strong policymaker $\{a_t(h_t, \tau = 1)\}_{t=0}^{\infty}$ and on the conjectured strategy profile of the weak policymaker $\{\bar{a}_t(h_t, \tau = 2)\}_{t=0}^{\infty}$.

At each period t, for any announced strategy profile of the strong policymaker and any conjectured strategy profile of the weak policymaker $\{a_1, \bar{a}_2\} \equiv \{a_t(h_t, \tau = 1), \bar{a}_t(h_t, \tau = 2)\}_{t=0}^{\infty}$, private agents' belief of a type $\tau = 1$ policymaker being present is *consistent* in the sense that it is formed according to Bayes' rule, conditional on their information set $\{h_t, \pi_t\}$:

$$\mathcal{R}_t(1, h_t, \pi_t | a_1, \bar{a}_2) \equiv \rho_{t+1} = \left\{ \frac{\rho_t \theta(\pi_t, a_t(h_t, \tau = 1))}{\rho_t \theta(\pi_t, a_t(h_t, \tau = 1)) + (1 - \rho_t)\theta(\pi_t, \bar{a}_t(h_t, \tau = 2))} \right\} \tag{10}$$

with initial value ρ_0. This consistent belief implies the discounted expectation of future policy outcome is:

$$\begin{aligned} \mathcal{E}_t(h_t, \pi_t | a_1, \bar{a}_2) &= \beta \mathcal{R}_t(1, h_t, \pi_t | a_1, \bar{a}_2) \sum_{\varsigma_{t+1} \in S} \delta(\varsigma_{t+1}, \varsigma_t) \sum_{\pi_{t+1} \in \Pi} \pi_{t+1} \theta(\pi_{t+1}, a_{t+1}(h_{t+1}, 1)) \\ &\quad + \mathcal{R}_t(2, h_t, \pi_t | a_1, \bar{a}_2) \sum_{\varsigma_{t+1} \in S} \delta(\varsigma_{t+1}, \varsigma_t) \sum_{\pi_{t+1} \in \Pi} \pi_{t+1} \theta(\pi_{t+1}, \bar{a}_{t+1}(h_{t+1}, 2)) \end{aligned}$$

where $\mathcal{R}_t(2, h_t, \pi_t | a_1, \bar{a}_2) = 1 - \rho_{t+1}$.

Only in equilibrium, we impose rational expectations on private agents so that the conjectured strategy profile coincide with the actual strategy profile implemented by the weak policymaker, $\{\bar{a}_t(h_t, \tau = 2)\}_{t=0}^{\infty} = \{a_t(h_t, \tau = 2)\}_{t=0}^{\infty}$. Thus, in equilibrium, the learning function $r_t(\tau, h_t, \pi_t)$ and the expectation function $e_t(h_t, \pi_t)$ should satisfy:

$$\begin{aligned} r_t(\tau, h_t, \pi_t) &= \mathcal{R}_t(\tau, h_t, \pi_t | a_1, a_2) \qquad (11) \\ e_t(h_t, \pi_t) &= \beta \sum_{\tau} \mathcal{R}_t(\tau, h_t, \pi_t | a_1, a_2) \sum_{\varsigma_{t+1} \in S} \delta(\varsigma_{t+1}, \varsigma_t) \sum_{\pi_{t+1} \in \Pi} \pi_{t+1} \theta(\pi_{t+1}, a_{t+1}(h_{t+1}, \tau)) \end{aligned}$$

3.2.3 Sequential rationality of the weak policymaker

Another implication of the "imperfect public monitoring" is that the weak policymaker does not have strategic power on the expectation formation of private agents, because it is impossible for private agents to adjust their learning based on some unobservable change of weak type's policy action.

As a result, although the weak policymaker acts prior to the private agents in each period, they are essentially playing a *simultaneous game*[10] – the weak policymaker takes as given the private agents' learning function $r_t(\tau, h_t, \pi_t)$ when he decides the policy action in period $t : a_t(h_t, 2)$; private agents take as given the policy action of the weak type at period t, $\bar{a}_t(h_t, 2)$, when they update their belief about policymaker type: $r_t(\tau, h_t, \pi_t)$.

Once the learning function $r_t(\tau, h_t, \pi_t)$ is taken as given, so is the expectation function $e_t(h_t, \pi_t)$, since the period t weak policymaker does not have control either on the future strategies of the strong policymaker, $\{a_s(h_s, \tau = 1)\}_{s=t+1}^{\infty}$, or on private agents' conjectured strategies of future weak policymakers, $\{\bar{a}_s(h_s, \tau = 2)\}_{s=t+1}^{\infty}$. Similarly, all future expectation functions $\{e_s(h_s, \pi_s)\}_{s=t+1}^{\infty}$ are exogenous to the period t weak policymaker.

At each period t, the strategy of the type $\tau = 2$ policymaker is *sequentially rational*, given any current and future expectation functions $\{e_s(h_s, \pi_s)\}_{s=t}^{\infty}$ and his own future strategies $\{a_s(h_s, \tau = 2)\}_{s=t+1}^{\infty}$:

$$
\begin{aligned}
& \mathcal{A}_t\left(h_t, \tau = 2 | \{e_s(h_s, \pi_s)\}_{s=t}^{\infty}, \{a_s(h_s, \tau = 2)\}_{s=t+1}^{\infty}\right) \\
= \quad & \arg\max_a E_t \widetilde{U}_t = E_t\{\sum_{s=t}^{\infty} b^{t-s} \tilde{u}(\pi_s, e_s(h_s, \pi_s), \varsigma_s)\}
\end{aligned}
$$

where $\pi_s = \pi$ with probability $\theta(\pi, a_s(h_s, \tau = 2))$.

3.2.4 Sequential rationality of the strong policymaker

By contrast, since the strong policymaker decides all his future actions at period 0, so he is the Stackelberg leader with respect to both private agents and the weak policymaker (type $\tau = 2$). Thus, the strategy profile of the strong policymaker is *sequentially rational* in the sense that it maximizes his expected payoff at period 0, $E_0 U_0$, taking into account the impact of his strategy profile on the expectation functions of private agents and on the strategies of the weak policy-

[10] The situation is analogous to Orborne and Rubinstein's [1991, pp. xxx] discussion of the Stackelberg game.

maker, with the understanding that private agents' expectations are formed based on a *consistent* belief system and are conditional on the equilibrium strategies of the weak policymaker (rational expectation), and the strategies of the weak policymaker are *sequentially rational.*

3.2.5 Definition

Definition 1 $\{a_t^*(h_t,\tau=1), a_t^*(h_t,\tau=2), e_t^*(h_t,\pi_t)\}_{t=0}^{\infty}$ *forms a Perfect Bayesian Public Equilibrium if*

i) At each period t, rational expectations must hold:

$$\mathcal{E}_t^*(h_t,\pi_t|a_1) = \mathcal{E}_t(h_t,\pi_t|a_1,\{\mathcal{A}_t^*(h_t,\tau=2|a_1)\}_{t=0}^{\infty})$$

where the strategy profile of the type $\tau = 2$ policymaker, $\{\mathcal{A}_t^*(h_t,\tau=2|a_1)\}_{t=0}^{\infty}$, is sequentially rational:

$$\mathcal{A}_t^*(h_t,\tau=2|a_1) = \mathcal{A}_t\left(h_t,\tau=2|\,\{\mathcal{E}_s^*(h_s,\pi_s|a_1)\}_{s=t}^{\infty}, \{\mathcal{A}_s^*(h_s,\tau=2|a_1)\}_{s=t+1}^{\infty}\right)$$

ii) At $t=0$, the strategy profile of the type $\tau=1$ policymaker is sequentially rational, given that the expectation function is obtained according to i):

$$\{a_t^*(h_t,\tau=1)\}_{t=0}^{\infty} = \arg\max_{\{a_t(\tau=1)\}_{t=0}^{\infty}} E_0U_0 = E_0\{\sum_{t=0}^{\infty}\beta^t u(\pi_t,\mathcal{E}_t^*(h_t,\pi_t|a_1),\varsigma_t)\}$$

where $\pi_t = \pi$ with probability $\theta(\pi, a_t^*(h_t,\tau=1))$.

iii) The equilibrium expectation function and the strategy profile of the weak policymaker are:

$$\begin{aligned} e_t^*(h_t,\pi_t) &= \mathcal{E}_t^*(h_t,\pi_t|\,\{a_t^*(h_t,\tau=1)\}_{t=0}^{\infty}) \\ a_t^*(h_t,\tau=2) &= \mathcal{A}_t^*(h_t,\tau=2|\,\{a_t^*(h_t,\tau=1)\}_{t=0}^{\infty}) \end{aligned}$$

3.2.6 The FOC of the weak policymaker

As argued above, the current and future expectation functions $\{e_s(h_s,\pi_s)\}_{s=t}^{\infty}$ are exogenous to the period t weak policymaker. In addition, the period t weak type treats his own future actions as exogenous functions of future history, $\{a_s(h_s,\tau=2)\}_{s=t+1}^{\infty}$. These functions effectively determine the continuation value function of the period t weak type as defined below:

$$V_{t+1}\left(h_{t+1}\right) \equiv E_{t+1} \sum_{s=t+1}^{\infty} b^{s-(t+1)} \widetilde{u}(\pi_s, e_s\left(h_s, \pi_s\right), \varsigma_s)$$

where $\pi_s = \pi$ with probability $\theta\left(\pi, a_s\left(h_s, \tau = 2\right)\right)$.

In sum, when the period t weak type chooses his policy action $a_t(h_t, 2)$, he takes as given private agents' outcome-contingent expectation function $e_t\left(h_t, \pi_t\right)$ and his continuation value function $V_{t+1}\left(h_{t+1}\right)$:

$$\begin{aligned} a_t(h_t, 2) &= \arg\max_a \{ \sum_{\pi_t \in \Pi} \theta(\pi_t, a) \widetilde{u}(\pi_t, e_t\left(h_t, \pi_t\right), \varsigma_t) \\ &\quad + b \sum_{\pi_t \in \Pi} \sum_{\varsigma_{t+1} \in S} \theta(\pi_t, a) \delta(\varsigma_{t+1}, \varsigma_t) V_{t+1}\left(h_{t+1} = (h_t, \pi_t, \varsigma_{t+1})\right), \end{aligned} \tag{12}$$

but recognizes that his current action a affects the likelihood for various outcomes π and different policy outcomes are part of the future histories that will affect both his momentary objective and his future value.

Hence, given $e_t\left(h_t, \pi_t\right)$ and $V_{t+1}\left(h_{t+1}\right)$, the period t weak type's policy action $a_t(h_t, 2)$ compatible with sequential rationality needs to satisfy the following first order condition:

$$\begin{aligned} 0 &= \{ \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_t(h_t, 2)) \widetilde{u}(\pi_t, e_t\left(h_t, \pi_t\right), \varsigma_t) \\ &\quad + b \sum_{\pi_t \in \Pi} \sum_{\varsigma_{t+1} \in S} \theta_a(\pi_t, a_t(h_t, 2)) \delta(\varsigma_{t+1}, \varsigma_t) V_{t+1}\left(h_{t+1} = (h_t, \pi_t, \varsigma_{t+1})\right) \end{aligned} \tag{13}$$

3.2.7 Value for the weak type

Given his current action $a_t(h_t, 2)$, the outcome-contingent expectation function $e_t\left(h_t, \pi_t\right)$ and his future value $V_{t+1}\left(h_{t+1}\right)$, the value of a period t weak type must satisfy:

$$\begin{aligned} V_t(h_t) &= \sum_{\pi_t \in \Pi} \theta(\pi_t, a_t(h_t, 2)) \widetilde{u}(\pi_t, e_t\left(h_t, \pi_t\right), \varsigma_t) \\ &\quad + b \sum_{\pi_t \in \Pi} \sum_{\varsigma_{t+1} \in S} \theta(\pi_t, a_t(h_t, 2)) \delta(\varsigma_{t+1}, \varsigma_t) V_{t+1}\left(h_{t+1} = (h_t, \pi_t, \varsigma_{t+1})\right) \end{aligned} \tag{14}$$

4 The Principal-Agent Approach

We aim to characterize and compute this equilibrium by studying a programming problem in which the type $\tau = 1$ policymaker maximizes his period 0 objective, taking into account *rational expectations* for the private sector and *sequential rationality* for all type $\tau = 2$ policymakers. In effect, as in the literature on the design of Ramsey plans, we view the strong policymaker as the principal which sets plans not only for his own future actions, but also for private agents' expectations and the weak type's future actions, prior to the play. When the strong type chooses private agents' expectations and the weak type's actions, it realizes that private agents' expectation is formed in a rational way and the weak type's action has to obey his sequential rationality. Our approach also thus relates to the design of optimal contracts along principal-agent lines, as we further elaborate below.

4.1 The strong's type policy problem

From the analysis above, the sequential rationality of a period t weak type and the rational expectation of private agents jointly restrict $a_t(h_t, 2)$ and $e_t(h_t, \pi_t)$, which must satisfy equations (13) and (11), given the actions of the strong type and the weak type's own future behavior. Moreover, the choices of policy actions $a_t(h_t, \tau = 1, 2)$ and expectations $e_t(h_t, \pi_t)$ govern the value of a period t weak type in history h_t, $V_t(h_t)$, according to the recursive specification (14).

Note that the strong type is assumed to choose the actions and value for the weak type, subject to the weak type's first order condition. In this regard, our work utilizes the insights of the "first order condition approach" to design of an optimal contracts. There must be some restrictions on the objective, $\widetilde{u}$, and the probability distribution, θ, for this approach to be valid: a future draft of this paper will spell these out in more detail. However, as in the static literature on principal-agent problems, the FOC approach allows for great simplification in the computation of an optimal policy.[11]

[11] Add references to static and dynamic optimal contract literature.

4.2 A dynamic Lagrangian

Since equations (11), (13) and (14) jointly restrict the choice set of the strong type, we can use these to form a Lagrangian that is suitable for determining the constrained optimal series of actions for the strong type: we perform the detailed construction in appendix A, but outline the key elements here. The "dynamic Lagrangian" takes the form,

$$L(h_0) = E_0 U_0 + \Psi_0^1 + \Psi_0^2 + \Psi_0^3$$

In this expression, $E_0 U_0$ is the strong policymaker's expected present discounted value objective. Ψ_0^1 is the Lagrangian component for the period-by-period rational expectation equations (11); Ψ_0^2 is the Lagrangian component for the sequential rationality of the weak type summarized by the FOCs (13); and Ψ_0^3 is the Lagrangian component for the value recursion (14). Each of the Ψ^i terms in the expression can be viewed as an expected present discounted value. We repeat the three forward-looking constraints, indicating the multiplier that we attach to it:

$$\begin{aligned}
\gamma_t(h_t, \pi_t) \quad &: \quad 0 = \beta \sum_{\tau} r_t(\tau, h_t, \pi_t) \sum_{\varsigma_{t+1} \in S} \delta(\varsigma_{t+1}, \varsigma_t) \mu(a_{t+1}(h_{t+1}, \tau)) - e_t(h_t, \pi_t) \\
\phi_t(h_t) \quad &: \quad 0 = \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_t(h_t, 2)) \widetilde{u}(\pi_t, e_t(h_t, \pi_t), \varsigma_t) \\
&\qquad + b \sum_{\pi_t \in \Pi} \sum_{\varsigma_{t+1} \in S} \theta_a(\pi_t, a_t(h_t, 2)) \delta(\varsigma_{t+1}, \varsigma_t) V_{t+1}(h_t, \pi_t, \varsigma_{t+1}) \\
\lambda_t(h_t) \quad &: \quad 0 = \sum_{\pi_t \in \Pi} \theta(\pi_t, a_t(h_t, 2)) \widetilde{u}(\pi_t, e_t(h_t, \pi_t), \varsigma_t) \\
&\qquad + b \sum_{\pi_t \in \Pi} \sum_{\varsigma_{t+1} \in S} \theta(\pi_t, a_t(h_t, 2)) \delta(\varsigma_{t+1}, \varsigma_t) V_{t+1}(h_t, \pi_t, \varsigma_{t+1}) - V_t(h_t)
\end{aligned}$$

Note that these "commitment multipliers" depend on the juncture at which the constraint is binding, which in turn depends on the sequencing of the actions described above. Expectations are formed after private agents see π_t, so that the multiplier reflects this dependence in addition to the history up to the start of the period, h_t. Weak type actions are taken and value is determined at the start of the period, so that the multiplier depends only on h_t.

There are three comments to be made about the dynamic Lagrangian. First, together with the transition equations, it is the conceptual basis for solving for history-dependent optimal policies under imperfect credibility from specified initial conditions. Second, as we show in appendix A, it can be written recursively in the style of Marcet and Marimon (2010), so that the methods of dynamic programming can be used to characterize and compute the optimal policy. Third, the dynamic programming problem determine plans which are made conditional on the strong type knowing that it is in power, so that the relevant conditional expectation for the objective is based on this knowledge. These plans are history dependent, but the effects of history are captured by a small number of state variables, which are explicable in their form and computationally desirable in their presence.

4.3 Recursive optimal policy program

As shown in appendix A, a small number of state variables are sufficient to summarize the history h_t. which significantly simplifies the optimal policy problem for the strong policymaker.

The recursive representation of the optimal policy problem for the type $\tau = 1$ policymaker is

$$\begin{aligned} & W(\varsigma_t, \rho_t, \eta_t) \\ = & \min_{\gamma_t(\pi_t), \phi_t, \lambda_t} \max_{a_{\tau t}, e_t(\pi_t), V_t} \{ w_t + \beta E_t W(\varsigma_{t+1}, \rho_{t+1}, \eta_{t+1}) \} \end{aligned} \tag{15}$$

where η_t is a three element vector of "pseudo-state variables", which evolve according to

$$\begin{aligned} \eta_{\gamma,t+1} &\equiv \gamma_t(\pi_t) \\ \eta_{\phi,t+1} &= \frac{b}{\beta} \phi_t \frac{\theta_a(\pi_t, a_{2t})}{\theta(\pi_t, a_{1t})} \\ \eta_{\lambda,t+1} &\equiv \frac{b}{\beta} \lambda_t \frac{\theta(\pi_t, a_{2t})}{\theta(\pi_t, a_{1t})} \end{aligned}$$

with the learning state variable evolving according to the Bayesian specification (10),

$$\rho_{t+1} = \left\{ \frac{\rho_t \theta(\pi_t, a_{1t})}{\rho_t \theta(\pi_t, a_{1t}) + (1 - \rho_t) \theta(\pi_t, a_{2t})} \right\}$$

and the exogenous shock ς_t evolving according to the specified Markov process (3).

We elaborate on the flow objective w_t. First, it contains expected utility for the strong type:

$$\sum_{\pi_t \in \Pi} \theta(\pi_t, a_{1t}) u(\pi_t, e_t(\pi_t), \varsigma_t)$$

Second, it contains current period variables and the pseudo-state variables that are constructed to be parsimonious combinations of lagged commitment multipliers and lagged variables.

$$\begin{aligned}
&[\eta_{\gamma,t}\rho_t\mu(a_{1t}) + \eta_{\gamma,t}(1-\rho_t)\mu(a_{2t}) - \sum_{\pi_t \in \Pi} \theta(\pi_t, a_{1t}) \gamma_t(\pi_t) e_t(\pi_t)] \\
&+[\phi_t \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_{2t}) \widetilde{u}(\pi_t, e_t(\pi_t), \varsigma_t) + \eta_{\phi,t} V_t] \\
&+[\lambda_t \sum_{\pi_t \in \Pi} \theta(\pi_t, a_{2t}) \widetilde{u}(\pi_t, e_t(\pi_t), \varsigma_t) - \lambda_t V_t + \eta_{\lambda,t} V_t]
\end{aligned}$$

Note that the programming problem specifies some choices – decisions (a_τ, V) and multipliers (ϕ, λ) – that are functions only of the state vector $\varsigma_t, \rho_t, \eta_t$ with the decisions (e) and multiplier (γ) also being functions of π. These latter "contingency plans" are a consequence of the fact that the private sector expectations are formed on the basis of more information than policymaker decisions, but expectations management is still a core feature of this environment.

5 Refining the equilibrium in announcements

Now let us revisit the announcement game occurred prior to the game of policy actions.

At the beginning of period $t = 0$, the type τ of the policymaker is drawn and then the policymaker in place makes its announcements $\alpha = \{\alpha_t(h_t, \tau)\}_{t=0}^{\infty}$.[12]

With probability g, it is the type $\tau = 1$ policymaker being drawn and he is trustworthy in the sense that his future policy action will follow the announcement. With probability $1-g$, it is the type $\tau = 2$ policymaker being drawn and he is opportunistic in the sense that he is free to deviate from the announced policy action in the actual implementation.

After seeing the announcements, the private agents will update their belief about the type of

[12] In the previous section, we show that the public history, $h_t = \{h_{t-1}, \pi_{t-1}, \varsigma_t\}$, is summarized by 5 state variables: $(\varsigma_t, \rho_t, \eta_{\gamma t}, \eta_{\phi t}, \eta_{\lambda t})$, which makes it feasible to announce a state-contingent policy rule depending on the all possible realized public history. In practice, the state variables $(\rho_t, \eta_{\gamma t}, \eta_{\phi t}, \eta_{\lambda t})$ are outcomes of $\{\pi_s\}_{s=0}^{t}$ so that they are observable to the private agents.

current policymaker, from g to ρ_0, which will in turn affect the formation of all their future **rational** expectations, i.e. $\{\mathcal{E}_t^*(h_t, \pi_t|\alpha)\}_{t=0}^{\infty}$ as defined in Definition 1, through Bayesian learning (10).

The relevant payoff for the type $\tau = 1$ sender is the one at period 0:

$$E(U_0|H_0) = E_0\{\sum_{t=0}^{\infty} \beta^t u(\pi_t, \mathcal{E}_t^*(h_t, \pi_t|\alpha), \varsigma_t)\}$$

where $\pi_t = \pi$ with probability $\theta(\pi, a_t(h_t, \tau = 1)) = \theta(\pi, \alpha_t(h_t, \tau = 1))$.

The relevant payoff for the type $\tau = 2$ sender is the one at period 0:

$$E(\widetilde{U}_0|\widetilde{H}_0) = E_0\{\sum_{t=0}^{\infty} b^t \widetilde{u}(\pi_t, \mathcal{E}_t^*(h_t, \pi_t|\alpha), \varsigma_t)\}$$

where $\pi_t = \pi$ with probability $\theta(\pi, a_t(h_t, \tau = 2))$.

We now solve for the signalling equilibria in this period-0 game.

5.1 The equilibrium is always pooling

If not, suppose $\{\alpha_t(h_t, \tau = 1)\}_{t=0}^{\infty} \neq \{\alpha_t(h_t, \tau = 2)\}_{t=0}^{\infty}$. By announcing $\{\alpha_t(h_t, \tau = 2)\}_{t=0}^{\infty}$, the type $\tau = 2$ policymaker will be identified and $\rho_0 = 0$, which implies all future $\rho_t = 0$. In addition,

$$E_0\{\sum_{t=0}^{\infty} b^t \widetilde{u}(\pi_t, \mathcal{E}_t^*(\rho_t = 0), \varsigma_t)\} < E_0\{\sum_{t=0}^{\infty} b^t \widetilde{u}(\pi_t, \mathcal{E}_t^*(\rho_t > 0), \varsigma_t)\}$$

so that the type $\tau = 2$ policymaker will be strictly better off by deviating to make the same announcements as the type $\tau = 1$ policymaker.

5.2 Multiple pooling equilibria

In fact, without any restriction on the off-equilibrium belief, any $\{\bar{\alpha}_t(h_t)\}_{t=0}^{\infty}$ can be the equilibrium message if $\rho_0 = 0$ for any announcement $\{\alpha_t(h_t)\}_{t=0}^{\infty} \neq \{\bar{\alpha}_t(h_t)\}_{t=0}^{\infty}$. Because $\rho_0 = 0$ gives the lowest possible payoff for both types of policymakers, this off-equilibrium threat by private agents will force both types of policymakers to announce $\{\bar{\alpha}_t(h_t)\}_{t=0}^{\infty}$ in equilibrium and the equilibrium belief is $\rho_0 = g$.

5.3 Refinement in the spirit of Mailath et. al. (1993)

We now apply a refinement in the spirit of Mailath et. al. (1993) to select a unique equilibrium out of the set of pooling equilibria. In Mailath et. al. (1993), they postulate three key ingredients in constructing belief restrictions in a candidate equilibrium. First, any out-of-equilibrium message must be one that is sent in an alternative equilibrium by some set of agents. Second, the incentives that various types of agents have to send an alternative message τ are evaluated by comparison of their benefits in a candidate and an alternative equilibrium. Third, when such comparison induces an out-of-equilibrium probability distribution of sender types computed by Bayes' law, labeled in this paper as the "strongly coherent belief", it is then used to generate restrictions on beliefs.

The equilibrium strategy profile of the type $\tau = 1$ policymaker obtained in Definition 1 is the one solves the following maximization problem by assuming $\rho_0 = g$:

$$\{a_t^* (h_t, \tau = 1)\}_{t=0}^{\infty} = \arg \max_{\{a_t(\tau=1)\}_{t=0}^{\infty}} E_0 U_0 = E_0\{\sum_{t=0}^{\infty} \beta^t u(\pi_t, \mathcal{E}_t^* (h_t, \pi_t|a_1), \varsigma_t)\}$$

where $\pi_t = \pi$ with probability $\theta (\pi, a_t^* (h_t, \tau = 1))$.

We will show that $\{a_t^* (h_t, \tau = 1)\}_{t=0}^{\infty}$ is the only equilibrium message that survives the refinement.

We first treat the pooling equilibrium with $\alpha^* = \{a_t^* (h_t, \tau = 1)\}_{t=0}^{\infty}$ as the candidate equilibrium with the equilibrium payoffs

$$R^* (\tau = 1) = W \left(\varsigma_0, \rho_0, \eta_{\gamma 0}, \eta_{\phi 0}, \eta_{\lambda 0}\right); R^* (\tau = 2) = V \left(\varsigma_0, \rho_0, \eta_{\gamma 0}, \eta_{\phi 0}, \eta_{\lambda 0}\right),$$

and treat a pooling equilibrium with any other $\bar{\alpha} = \{\bar{\alpha}_t (h_t)\}_{t=0}^{\infty}$ as the alternative equilibrium with the equilibrium payoffs

$$\begin{aligned} \bar{R} (\tau = 1) &= E_0\{\sum_{t=0}^{\infty} \beta^t u(\pi_t, \mathcal{E}_t^* (h_t, \pi_t|\bar{\alpha}), \varsigma_t)\}; \\ \bar{R} (\tau = 2) &= E_0\{\sum_{t=0}^{\infty} \beta^t \tilde{u}(\pi_t, \mathcal{E}_t^* (h_t, \pi_t|\bar{\alpha}), \varsigma_t)\}. \end{aligned}$$

Because $R^* (\tau = 1) > \bar{R} (\tau = 1)$ by construction, only the type $\tau = 2$ policymaker may have incentive to deviate to message $\bar{\alpha}$. Thus, the strongly-coherent belief at message $\bar{\alpha}$, $\varphi (\tau = 1|\alpha = \bar{\alpha}) =$

0, which is consistent with the off-equilibrium belief $\Pr(\tau = 1|\bar{\alpha}) = 0$. The pooling equilibrium with $\alpha^* = \{a_t^*(h_t, \tau = 1)\}_{t=0}^{\infty}$ then survives the refinement.

Now we treat the pooling equilibrium with $\bar{\alpha} = \{\bar{\alpha}_t(h_t)\}_{t=0}^{\infty} \neq \alpha^*$ as the candidate equilibrium and the pooling equilibrium with $\alpha^* = \{a_t^*(h_t, \tau = 1)\}_{t=0}^{\infty}$ as the alternative equilibrium. We know that to support the equilibrium $\bar{\alpha} = \{\bar{\alpha}_t(h_t)\}_{t=0}^{\infty}$, we need $\Pr(\tau = 1|\alpha^*) < g = \Pr(\tau = 1|\bar{\alpha})$ to make it optimal for the type $\tau = 1$ policymaker to send the equilibrium message $\bar{\alpha}$. However, we know that the payoff in the alternative equilibrium is larger than the payoff in the candidate equilibrium for the type $\tau = 1$ policymaker ($R^*(\tau = 1) > \bar{R}(\tau = 1)$ by construction). Thus, the strongly-coherent belief at message α^*, $\varphi(\tau = 1|\alpha = \alpha^*)$, should be at least g :

$$\varphi(\tau = 1|\alpha = \alpha^*) \geqslant g,$$

which is inconsistent with the supporting belief

$$\Pr(\tau = 1|\alpha^*) < g.$$

Therefore, the pooling equilibrium with $\bar{\alpha} = \{\bar{\alpha}_t(h_t)\}_{t=0}^{\infty}$ is ruled out by the refinement.

6 Conclusion and Remarks

[TO BE ADDED]

References

[1] Aoyagi, Masaki. 1993. "Reputation and Dynamic Stackelberg Leadership in Infinitely Repeated Games." *Journal of Economic Theory*, 71:378-393.

[2] Backus, David A, and John Driffill. 1985a. "Inflation and Reputation." *American Economic Review*, 75(3): 530-538.

[3] 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.

[4] Ball, Laurence. 1995. "Time-consistent Policy and Persistent Changes in Inflation." *Journal of Monetary Economics*, 36: 329-350.

[5] Barro, Robert J. 1986. "Reputation in a Model of Monetary Policy with Incomplete Information." *Journal of Monetary Economics*, 17: 3-20.

[6] Barro, Robert J, and David B. Gordon. 1983. "Rules, discretion and reputation in a model of monetary policy." *Journal of Monetary Economics*, 12: 101-122.

[7] Chatterjee, Satyajit, Dean Corbae and Jose-Victor Rios-Rull. 2009. "A Theory of Credit Scoring and Competitive Pricing of Default Risk." http://pier.econ.upenn.edu/Events/scorevictor.pdf.

[8] Celentani, Marco, Drew Fudenberg, David K. Levine and Wolfgang Pesendorfer. 1996. "Maintaining a Reputation Against a Long-Lived Opponent." *Econometrica*, 64(3): 691-704.

[9] Cole, Harold L., James Dow, and William B. English. 1995. "Default, Settlement, and Signalling: Lending Resumption in a Reputational Model of Sovereign Debt." *International Economic Review*, 36: 365-385.

[10] Cole, Harold L. and Patric J. Kehoe. 1998. "Models of Sovereign Debt: Partial versus General Reputations." *International Economic Review*, 39: 55-70.

[11] Crawford, Vincent P. and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica*, 50(6): 1431-1451.

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

[13] Cukierman, Alex. 1993. *Central Bank Strategy, Credibility, and Independence.* Cambridge, MA: MIT Press.

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

[15] Cukierman, Alex, and Allan H. Meltzer. 1986. "A Theory of Ambiguity, Credibility, and Inflation under Asymmetric Information" *Econometrica*, 54: 1099-1128.

[16] Diamond, Douglas W. 1989. "Reputation Acquisition in Debt Markets." *The Journal of Political Economy*, 97(4): 828-862.

[17] Evans, Robert and Jonathan P. Thomas. 1997. "Reputation and Experimentation in Repeated Games with Two Long-Run Players." *Econometrica*, 65(5): 1153-1173.

[18] Fischer, Stanley. 1990. "Rules versus Discretion in Monetary Policy." in *Handbook of Monetary Economics*, ed. Benjamin M. Friedman and Frank H. Hahn, 1155-1184. Amsterdam: North-Holland.

[19] Fudenberg, Drew and David K. Levine. 1983. "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games." *Journal of Economic Theory*, 31: 251-268.

[20] Fudenberg, Drew and David K. Levine. 1989. "Reputation and Equilibrium Selection in Games with a Patient Player." *Econometrica*, 57(4): 759-778.

[21] Fudenberg, Drew and David K. Levine. 1992. "Maintaining a Reputation when Strategies are Imperfectly Observed." *The Review of Economic Studies*, 59(3): 561-579

[22] Ghosh, Sambuddha. 2012. "Multiple Opponents and The Limits of Reputation." http://people.bu.edu/sghosh/Academic_files/combined_ch1_6.5.10.pdf

[23] Grossman, Sanford J, and Motty Perry. 1986. "Perfect Sequential Equilibrium." *Journal of Economic Theory*, 39(1): 97-119.

[24] Horner, Johannes. 2002. "Reputation and Competition." *The American Economic Review*, 92(3): 644-663

[25] Ireland, Peter N., Comments on Stokey: http://fmwww.bc.edu/ec-p/wp530.pdf

[26] King, Robert G., Yang K. Lu and Ernesto Pasten 2008. "Managing Expectations." *Journal of Money, Credit and Banking*, 40(8): 1625-1666.

[27] Kreps, David M. and Robert Wilson. 1982a. "Sequential Equilibria." *Econometrica*, 50: 863-894.

[28] Kreps, David M. and Robert Wilson. 1982b. "Reputation and Imperfect Information." *Journal of Economic Theory*, 27: 253-279.

[29] Kydland, Finn E. and Edward C. Prescott. 1977. "Rules rather than discretion." *Journal of Political Economy*, 85: 473-491.

[30] Lu, Yang K. 2013. "Optimal Policy with Credibility Concerns" http://ihome.ust.hk/~yanglu/LU_CredibilityConcern.pdf

[31] Mailath, George J, Masahiro Okuno-Fujiwara and Andrew Postlewaite. 1993. "Belief-Based Refinements in Singnalling Games." *Journal of Economic Theory*, 60: 241-276.

[32] Mailath, George J. and Larry Samulson. 2001. "Who Wants a Good Reputation?" *Review of Economic Studies*, 68: 415-441.

[33] Maskin, Eric and Jean Tirole. 1992. "The Principal-Agent Relationship with an Informed Principal, II: Common Values." *Econometrica*, 60(1): 1-42.

[34] Maskin, Eric and Jean Tirole. 2001. "Markov Perfect Equilibrium: I. Observable Actions." *Journal of Economic Theory*, 100: 191-219.

[35] Milgrom, Paul and John Roberts. 1982. "Predation, Reputation, and Entry Deterrence." *Journal of Economic Theory*, 27: 280-312.

[36] Myerson, Roger B. 1983. ""Mechanism Design by an Informed Principal." *Econometrica*, 51(6): 1767-1797.

[37] Phelan, Chirstopher. 2006. "Public trust and Government Betrayal." *Journal of Economic Theory*, 130: 27-43.

[38] Schmidt, Klaus M. 1993. "Reputation and Equilibrium Characterization in Repeated Games with Conflicting Interests." *Econometrica*, 61(2): 325-351.

[39] Sobel, Joel. 1985. "A Theory of Credibility." *The Review of Economic Studies*, 52(4): 557-573.

[40] Stokey, Nancy L. 2002. ""Rules vs. Discretion" after Twenty Five Years." NBER Macro Annual, 17

Appendix A:

Derivation of recursive policy program for the strong policymaker

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

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

and the expected, present-discounted value of the period t weak type's momentary objective is:

$$\tilde{E}_t \tilde{U}_t = \sum_{s=t}^{\infty} \beta^{s-t} \sum_{h_s} p(h_s) \sum_{\pi_s \in \Pi} \theta(\pi_s, a_s(h_s, \tau = 2))\, \tilde{u}(\pi_s, e_s(h_s, \pi_s), \varsigma_s).$$

Notice that the probability assessments is conditional on the type of policymaker who is present for the game.

A.1 The Lagrangian component for the rational expectation constraints

The strong policymaker in each period takes policy action before private agents form their expectation, which is contingent on the realization of policy outcome π_t. So we attach multipliers $p(h_t) \sum_{\pi_t \in \Pi} \theta(\pi_t, a_t(h_t, \tau = 1))\, \gamma_t(h_t, \pi_t)$ to (11) and produce the Lagrangian component:

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

Rearranging this, we obtain

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

if we assume $\gamma_{-1}(h_{-1}, \pi_{-1}) = 0$. The first equality makes use of the fact that $p(h_{t+1}) = p(h_t)\,\theta(\pi_t, a_t(h_t, 1))\,\delta(\varsigma_{t+1}, \varsigma_t)$ 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 $e_t(\pi_t) = e_t(h_t, \pi_t)$;[13] 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

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

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).

A.2 The Lagrangian component for the sequential rationality of the weak type

We assume that the sequential rationality of a period t weak type can be summarized by the FOC (13), which is in line with the "First Order Condition Approach" in the literature on principal-agent problems. We can attach multipliers, $p(h_t)\,\phi_t(h_t)$, to these period-by-period constraints. The Lagrangian component for the weak type FOC is:

$$\begin{aligned}
\Psi_0^2 &= \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t)\,\phi_t(h_t) \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_t(h_t, \tau = 2))\widetilde{u}(\pi_t, e_t(h_t, \pi_t), \varsigma_t) \\
&\quad + \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t)\,\phi_t(h_t) \left[b \sum_{\pi_t \in \Pi} \sum_{\varsigma_{t+1} \in S} \theta_a(\pi_t, a_t(h_t, \tau = 2))\delta(\varsigma_{t+1}, \varsigma_t) V_{t+1}(h_{t+1} = (h_t, \pi_t, \varsigma_{t+1})) \right] \\
&= \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t)\,\phi_t(h_t) \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_t(h_t, \tau = 2))\widetilde{u}(\pi_t, e_t(h_t, \pi_t), \varsigma_t) \\
&\quad + \frac{b}{\beta} \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t)\,\phi_{t-1}(h_{t-1}) \frac{\theta_a(\pi_{t-1}, a_{t-1}(h_{t-1}, \tau = 2))}{\theta(\pi_{t-1}, a_{t-1}(h_{t-1}, \tau = 1))} V_t(h_t) \\
&= E_0 \sum_{t=0}^{\infty} \beta^t \{\phi_t \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_{2t})\widetilde{u}(\pi_t, e_t(\pi_t), \varsigma_t) + \frac{b}{\beta}\phi_{t-1} \frac{\theta_a(\pi_{t-1}, a_{2,t-1})}{\theta(\pi_{t-1}, a_{1,t-1})} V_t\}
\end{aligned}$$

[13] 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.

if we assume $\phi_{-1}(h_{-1}) = 0$. The short-hand notation used here is $\phi_t = \phi_t(h_t)$ and $V_t(h_t) = V_t$. Notice that there is a change of measure in each of these constraints, reflecting the fact that the probability assessments that are relevant for the multipliers are those of the strong type.

The recursive form for this Lagrangian component is thus

$$\Psi_t^2 = \left[\phi_t \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_{2t}) \widetilde{u}(\pi_t, e_t(\pi_t), \varsigma_t) + \frac{b}{\beta} \phi_{t-1} \frac{\theta_a(\pi_{t-1}, a_{2,t-1})}{\theta(\pi_{t-1}, a_{1,t-1})} V_t\right] + \beta E_t \Psi_{t+1}^2$$

A.3 The Lagrangian component for the value recursion

Similarly, the Lagrangian for the Bellman equations (14) is, after attaching multipliers $p(h_t)\lambda_t(h_t)$ to each one of them in period t:

$$\begin{aligned}
\Psi_0^3 &= \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t) \lambda_t(h_t) \left[\sum_{\pi_t \in \Pi} \theta(\pi_t, a_t(h_t, \tau = 2)) \widetilde{u}(\pi_t, e_t(h_t, \pi_t), \varsigma_t) - V_t(h_t)\right] \\
&\quad + \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t) \lambda_t(h_t) \left[b \sum_{\pi_t \in \Pi} \sum_{\varsigma_{t+1} \in S} \theta(\pi_t, a_t(h_t, \tau = 2)) \delta(\varsigma_{t+1}, \varsigma_t) V_{t+1}(h_{t+1} = (h_t, \pi_t, \varsigma_{t+1}))\right] \\
&= \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t) \lambda_t(h_t) \left[\sum_{\pi_t \in \Pi} \theta(\pi_t, a_t(h_t, \tau = 2)) \widetilde{u}(\pi_t, e_t(h_t, \pi_t), \varsigma_t) - V_t(h_t)\right] \\
&\quad + \frac{b}{\beta} \sum_{t=0}^{\infty} \beta^t \sum_{h_t} p(h_t) \lambda_{t-1}(h_{t-1}) \frac{\theta(\pi_{t-1}, a_{t-1}(h_{t-1}, \tau = 2))}{\theta(\pi_{t-1}, a_{t-1}(h_{t-1}, \tau = 1))} V_t(h_t). \\
&= E_0 \sum_{t=0}^{\infty} \beta^t \{\lambda_t [\sum_{\pi_t \in \Pi} \theta(\pi_t, a_{2t}) \widetilde{u}(\pi_t, e_t(\pi_t), \varsigma_t) - V_t] + \frac{b}{\beta} \lambda_{t-1} \frac{\theta(\pi_{t-1}, a_{2,t-1})}{\theta(\pi_{t-1}, a_{1,t-1})} V_t\}
\end{aligned}$$

with $\lambda_t = \lambda_t(h_t)$ and λ_{-1} assumed to be zero. Notice again the change of measure here.

In turn, the recursive form for this Lagrangian component is:

$$\Psi_t^3 = \left[\lambda_t \sum_{\pi_t \in \Pi} \theta(\pi_t, a_{2t}) \widetilde{u}(\pi_t, e_t(\pi_t), \varsigma_t) - \lambda_t V_t + \frac{b}{\beta} \lambda_{t-1} \frac{\theta(\pi_{t-1}, a_{2,t-1})}{\theta(\pi_{t-1}, a_{1,t-1})} V_t\right] + \beta E_t \Psi_{t+1}^3$$

A.4 Recursive optimal policy program

Up to now, we have obtained explicit forms for all the components in the "dynamic Lagrangian" defined in section 4.2:

$$L(h_0) = E_0 U_0 + \Psi_0^1 + \Psi_0^2 + \Psi_0^3.$$

In addition, each component can be expressed in a recursive form:

$$E_t U_t = \sum_{\pi_t \in \Pi} \theta\left(\pi_t, a_{1t}\right) u(\pi_t, e_t\left(\pi_t\right), \varsigma_t) + \beta E_t U_{t+1}$$

$$\begin{aligned}
\Psi_t^1 &= \left[\eta_{\gamma,t}\rho_t \mu(a_{1t}) + \eta_{\gamma,t}(1-\rho_t)\mu(a_{2t}) - \sum_{\pi_t \in \Pi} \theta\left(\pi_t, a_{1t}\right) \gamma_t\left(\pi_t\right) e_t(\pi_t) \right] + \beta E_t \Psi_{t+1}^1 \\
\Psi_t^2 &= \left[\phi_t \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_{2t}) \widetilde{u}(\pi_t, e_t\left(\pi_t\right), \varsigma_t) + \eta_{\phi,t} V_t \right] + \beta E_t \Psi_{t+1}^2 \\
\Psi_t^3 &= \left[\lambda_t \sum_{\pi_t \in \Pi} \theta_a(\pi_t, a_{2t}) \widetilde{u}(\pi_t, e_t\left(\pi_t\right), \varsigma_t) - \lambda_t V_t + \eta_{\lambda,t} V_t \right] + \beta E_t \Psi_{t+1}^3
\end{aligned}$$

where we simplify the notation by defining the following "pseudo state variables" in period t that are parsimonious combinations of lagged commitment multipliers and lagged variables:

$$\begin{aligned}
\eta_{\gamma,t+1} &\equiv \gamma_t\left(\pi_t\right) \\
\eta_{\phi,t+1} &= \frac{b}{\beta} \phi_t \frac{\theta_a(\pi_t, a_{2t})}{\theta(\pi_t, a_{1t})} \\
\eta_{\lambda,t+1} &\equiv \frac{b}{\beta} \lambda_t \frac{\theta(\pi_t, a_{2t})}{\theta(\pi_t, a_{1t})}
\end{aligned}$$

All together, we can place the optimal policy problem for the type $\tau = 1$ policymaker in a recursive form:

$$\begin{aligned}
& W(\varsigma_t, \rho_t, \eta_t) \\
= & \min_{\gamma_t(\pi_t), \phi_t, \lambda_t} \max_{a_{\tau t}, e_t(\pi_t), V_t} \left\{ w_t + \beta E_t W(\varsigma_{t+1}, \rho_{t+1}, \eta_{t+1}) \right\}
\end{aligned}$$

where η_t is the three element vector described above.

The flow objective w_t contains expected utility for the strong type:

$$\sum_{\pi_t \in \Pi} \theta\left(\pi_t, a_{1t}\right) u(\pi_t, e_t\left(\pi_t\right), \varsigma_t)$$

and all of the bracketed terms from the Lagrangian components, i.e.,

$$\begin{aligned}
&[\eta_{\gamma,t}\rho_t\mu(a_{1t}) + \eta_{\gamma,t}(1-\rho_t)\mu(a_{2t}) - \sum_{\pi_t\in\Pi}\theta\left(\pi_t, a_{1t}\right)\gamma_t\left(\pi_t\right)e_t(\pi_t)] \\
&+[\phi_t\sum_{\pi_t\in\Pi}\theta_a(\pi_t, a_{2t})\widetilde{u}(\pi_t, e_t\left(\pi_t\right), \varsigma_t) + \eta_{\phi,t}V_t] \\
&+[\lambda_t\sum_{\pi_t\in\Pi}\theta_a(\pi_t, a_{2t})\widetilde{u}(\pi_t, e_t\left(\pi_t\right), \varsigma_t) - \lambda_t V_t + \eta_{\lambda,t}V_t]
\end{aligned}$$

This recursive formulation highlights that the period t values of these psuedo-state variables η_t are sufficient to summarize the state of the system in t for the strong policymaker, when coupled with private agents' belief about type in period t, ρ_t, evolving over time according to Bayes' rule:

$$\rho_{t+1} \equiv \left\{\frac{\rho_t\theta(\pi_t, a_{1t})}{\rho_t\theta(\pi_t, a_{1t}) + (1-\rho_t)\theta(\pi_t, a_{2t})}\right\}$$

and the exogenous shock ς_t governed by the Markov process:

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