Rules without Commitment: Reputation and Incentives*

Alessandro Dovis
University of Pennsylvania
and NBER
adovis@upenn.edu

Rishabh Kirpalani
University of Wisconsin-Madison
rishabh.kirpalani@wisc.edu

September 2020

Abstract

This paper studies the optimal design of rules in a dynamic model when there is a time inconsistency problem and uncertainty about whether the policy maker can commit to follow the rule ex post. The policy maker can either be a commitment type, which can always commit to follow rules, or an optimizing type, which sequentially decides whether to follow rules or not. This type is unobservable to private agents, who learn about it through the actions of the policy maker. Higher beliefs that the policy maker is the commitment type (i.e., the policy maker's reputation) help promote good behavior by private agents. We show that in a large class of economies, preserving uncertainty about the policy maker's type is preferable from an ex-ante perspective. If the initial reputation is not too high, the optimal rule is the strictest one that is incentive compatible for the optimizing type. We show that reputational considerations imply that the optimal rule is *more lenient* than the one that would arise in a static environment. Moreover, *opaque* rules are preferable to transparent ones if reputation is high enough.

First version: June 2019. We thank Mark Aguiar, Gadi Barlevy, V. V. Chari, Russ Cooper, John Geanakoplos, Ben Hébert, Juanpa Nicolini, Erik Madsen, Ramon Marimon, Giuseppe Moscarini, Guillermo Ordonez, Chris Phelan, Facundo Piguillem, Aleh Tsyvinski, Pierre Yared, and Fabrizio Zilibotti for valuable comments. We would also like to thank Veronica Guerrieri and three anonymous referees for very helpful suggestions.

1 Introduction

Since Kydland and Prescott (1977), a large literature in macroeconomics has grappled with the problem of designing policies when there are time inconsistency problems. Rules are often proposed as a solution to the time inconsistency problem. The implicit assumption is that society can credibly impose rules on policy makers and that policy makers can commit to follow these rules. However, at the time when rules and regulations are formulated, there is often substantial *uncertainty* about whether policy makers can resist the temptation to deviate ex post from the stated rules if it is optimal for them to do so. This uncertainty is only resolved over time as the actions of policy makers are observed. The combination of uncertainty and learning generates reputational incentives for policy makers.

The key question motivating this paper is, how should rules be designed taking into account both the uncertainty about the policy makers' ability to follow the rules ex post and their reputation-building incentives? To answer it we study the optimal design of policy rules in a dynamic game between policy makers and private agents in which the policy maker's ability to commit is private information. We define the public beliefs about the ability of the policy maker to commit as the policy maker's *reputation*. The main result of our paper is that if the initial reputation is low enough, the optimal rule should be designed to preserve uncertainty in future periods. This is implemented by introducing leniency in policy. In contrast, if the initial reputation is high, the optimal rule should promote learning about this type. We also show that designing opaque rules can be beneficial when reputation is high since they help preserve uncertainty without the need to introduce leniency in rules.

The insights from our theory can be applied to many relevant policy design questions including the design of central bank mandates, fiscal rules in federal governments, and financial regulation. Consider, for instance, the optimal design of financial regulation. As is well understood, in a large class of economies, if regulators can commit, a no-bailout policy is optimal in order to prevent excessive risk taking by financial institutions ex ante. In particular, creditors should be forced to take losses in the event of default (bail-in). If the reputation of regulators is not sufficiently high, our analysis suggests that allowing for partial bailouts in equilibrium is optimal. We show that, contrary to conventional wisdom, bailouts along the equilibrium path are necessary to discipline future risk-taking of financial firms, as they preserve uncertainty about the type of the policy maker.

We consider a dynamic model with three types of agents: a rule designer, policy makers, and private agents. The rule designer chooses a rule, which consists of a policy recommendation to policy makers, in order to maximize the expected social welfare. After the rule is chosen, the private agents take their actions and, finally, the policy maker chooses a policy. As in Barro (1986), the policy maker can be one of two types: a commitment type, which always follows the recommendation, or an optimizing type, which follows the recommendation only if it is sequentially optimal to do so. This type is unobservable to both the rule designer and private agents. We define the beliefs that the policy maker is the commitment type as its reputation.

We present two leading examples of our framework. The first is a model similar to Barro and Gordon (1983b) in which the rule designer must choose the optimal inflation target. The second is a banking model in the spirit of Kareken and Wallace (1978) where there is a trade-off between providing incentives to bankers for taking appropriate levels of risk ex ante and bailing them out ex post to avoid a costly default. In this case the rule designer chooses an optimal bailout policy.

We first study a static problem. We show that under certain conditions, *uncertainty is beneficial* in that the expected social welfare is higher when the private agents and the rule designer are uncertain about the type of the policy maker relative to the case in which this type is revealed right before the rule designer chooses the rule. That is, the rule designer's static value is concave in the policy maker's reputation.[1] There are two critical conditions that generate this result. First, society's preferences are concave, or, equivalently, there are convex costs of deviating from the Ramsey outcome. Second, the equilibrium private action and the policy maker's static best response are strategic complements and this complementarity is stronger when reputation is low. In the context of the bailout example, this second condition implies that an increase in reputation incentivizes banks to take on less risk, and that the disciplining effect of reputation is greater when reputation is low and banks are taking on a lot of risk. Together, these two conditions along with a technical condition, imply that there are decreasing returns to reputation and that the rule designer's value is concave in reputation.

We then consider a repeated version of this policy game. We assume that at the beginning of each period the rule designer can revise the rules without commitment.[2] Unlike in the static model, the optimizing type now cares about its reputation in the following period as it affects the actions of the private agents. Thus, it can be incentivized to choose policies other than its static best response. We show that when reputation is low, the rule designer wants to preserve uncertainty about the type of the policy maker. The optimal rule in this case is the most stringent policy that is incentive compatible for the optimizing type. This recommended policy is more lenient than the statically optimal one. Leniency

[1] Nosal and Ordoñez (2016) also consider an environment in which uncertainty can mitigate the time inconsistency problem. The mechanism is very different: here there is uncertainty about the policy maker's type, while in their paper there is uncertainty about the state of the economy, which restrains the policy maker ex post.

[2] In Appendix E, we show that our main results are unchanged if the rule designer can commit, for sufficiently high or low levels of reputation.

in the rule makes it easier for the optimizing type to follow the recommendation ex post. This has dynamic benefits because it prevents the private agents from learning the type of the policy maker, and uncertainty is beneficial. When reputation is low, inducing the optimizing type to follow the rule also has static benefits. This is because it promotes better behavior by the private agents who anticipate that the optimizing type will follow the rule – albeit more lenient – instead of the statically optimal policy.

This result has sharp implications for policy design. In the context of optimal inflation targeting, having looser inflation targets is beneficial when reputation is low. Another application of our framework is the design of exchange rate regimes. Our result suggests that when reputation is low, crawling pegs might be superior to fixed-exchange-rate policies. Similarly, in the context of financial regulation, if reputation is low, the optimal rule is not a strict no-bailout policy that imposes losses on lenders. By explicitly allowing for partial bailouts along the equilibrium path, the rule designer makes it easier for the optimizing type to adhere to the rule and maintain its reputation. The optimal rule prescribed by the model is in contrast with the observed design of financial regulation after the 2008 financial crisis. After the bailouts of financial institutions during this crisis, the reputation of regulators was arguably low. While our model prescribes a more lenient bailout policy in this situation, the Dodd–Frank Act imposed strict no-bailout policies.

In contrast, if reputation is sufficiently high, the rule designer finds it optimal to set stringent rules that result in the type of the policy maker being revealed. This is because when reputation is sufficiently high, there are static costs associated with choosing a lenient rule. In this case, the private agents anticipate that the rule will be followed with sufficiently high probability and so by choosing the Ramsey policy, the rule designer can obtain a value close to the Ramsey outcome. There are, however, dynamic losses associated with choosing the Ramsey policy: if the rule is to follow the Ramsey policy, for a low enough discount factor, the optimizing type will not follow the rule and there will be revelation about the type of the policy maker in the first period. Because uncertainty is beneficial, the expected continuation value is lower than in the case in which the type of the policy maker is not revealed. When reputation is high enough, the static benefits of choosing a stringent rule outweigh the dynamic losses.

Next, we study the optimal degree of transparency of the rule. We say that a rule is transparent if the policy maker's deviations are easily detectable. In repeated policy games with no reputational considerations, perfect monitoring is always desirable. See Atkeson and Kehoe (2001), Atkeson et al. (2007), and Piguillem and Schneider (2013). In contrast, we show that with reputational considerations, transparent rules are desirable only for low levels of reputation, while opaque rules are desirable for high levels of reputation.[3] This is because they can help maintain reputation without the static costs associated with pooling when reputation is high.

We consider two ways in which the rule designer can affect the transparency of the rules. First, we assume that future private agents and rule designers observe only a signal of the chosen policy and the rule designer can choose the precision of the signal. High precision (transparency) is beneficial because it incentivizes the optimizing type to follow the rule, as a deviation results in the revelation of its type with large reputation losses. Low precision (opaqueness) is beneficial because it allows the rule designer to maintain uncertainty about the policy maker's type. For instance, if the signals are imprecise, the private agents attribute the observed deviations from the stated policy to noise rather than to the policy maker being the optimizing type that deviated from the policy. This is helpful for high levels of reputation since the rule designer would like to choose the Ramsey policy from a static perspective. As discussed earlier, there is a trade-off between the static value of having the commitment type follow a stringent rule and the dynamic losses associated with learning the policy maker's type. Allowing for opaque rules helps break this trade-off: the rule designer can achieve the high static payoff of choosing a rule equal to the Ramsey policy without the costs associated with separation for sure because the policy observations are very noisy.

An alternative way of introducing opacity in rules is to allow the rule designer to choose stochastic rules even though fundamentals are deterministic. When reputation is low, the optimal rule has no randomization in order to maximize the incentives of the optimizing type to follow more stringent policies. When reputation is high instead, it is optimal to have randomization in order to reduce the dispersion in the posteriors.

In our baseline setup, we model the commitment type as a policy maker that cannot deviate from the rules. One interpretation of this is that the commitment type suffers a cost from deviating from the stated rule over and above the reputational cost in the model. For example, a deviation may affect the commitment's type ability to be elected to higher offices, while the optimizing type may not have such ambitions. Alternatively, one could assume that the policy makers are identical, but there is uncertainty about whether these policy deviations can be enacted, due to legislative holdups, for example. In particular, policy makers always have an incentive to choose policies which are sequentially rational but might face roadblocks in their implementation if the legislature is controlled by opponents who might block these policies for purely political purposes. As in Piguillem and Riboni (2018), the rule can be the default option in case of such disagreements. In this case, we can interpret the commitment type as a policy maker which faces such roadblocks, and the optimizing type as one which does not. The latter might want to pretend

[3] In the principal–agent literature there are examples of environments where imperfect monitoring is beneficial to provide incentives. See, for instance, Crémer (1995) and Prat (2005).

that its hands are tied (like the commitment type) for exactly the same reasons as in the baseline model.

Related literature The macroeconomics literature has followed two approaches to modeling time inconsistency in policy games. First, in the tradition of Kydland and Prescott (1977), there are models in which the policy maker is benevolent and time-inconsistency problems arise because of externalities associated with the presence of private agents (see Chari et al. (1988)). Examples include Barro and Gordon (1983b), Chari and Kehoe (1990), Athey et al. (2005), and Kareken and Wallace (1978). Our setup follows in this tradition. Second, the literature also considers economies in which the time-inconsistency problem arises from a conflict of interest between the society and the policy makers. In these economies private agents do not play a critical role. Examples include Amador et al. (2006), Amador and Bagwell (2013), and Halac and Yared (2014). The rule designer's desire to influence the actions of the private agents is critical for our results. We show that in models that follow the second approach and do not have private agents, uncertainty is never beneficial in that the rule designer always prefers to learn the policy maker's type.

This paper is related to the literature that studies the trade-off between rules and flexibility. See, for example, Athey et al. (2005), Halac and Yared (2014), Halac and Yared (2018a), and Azzimonti et al. (2016), among others. The focus of this literature is on how much flexibility to leave the policy maker when it is not possible to make the rule contingent on the state of the economy (say because it is private information to the policy maker). We abstract from this issue by considering a deterministic environment, but we focus instead on the uncertainty about the ability of the policy maker to commit. Our paper is also related to the literature that studies optimal policies without commitment when it is known that the policy maker cannot commit. This is the approach followed by a large literature on time-consistent policies, including Barro and Gordon (1983b), Chari and Kehoe (1990), Phelan and Stacchetti (2001), and Halac and Yared (2018b). Our paper nests simple versions of these two approaches as special cases when the policy maker's reputation is either one or zero.

This paper builds on the reputation literature that originates with Milgrom and Roberts (1982) and Kreps and Wilson (1982). See Barro (1986), Backus and Driffill (1985), Phelan (2006), Amador and Phelan (2018), and Dovis and Kirpalani (2020b) for recent applications to policy games. Most of this literature takes as given the policy chosen by the commitment type and analyses the incentives of the optimizing type and the outcomes that can be achieved. The goal of this paper is to study the optimal policy that the commitment type should follow.

A key driver of our results is the idea that uncertainty about the policy makers' type is beneficial. This feature is also present in Dovis and Kirpalani (2020a). Our contribution is to show how this property affects the design of the optimal rule. Marinovic and Szydlowski (2019), Bond and Zeng (2018), and Asriyan et al. (2019) also consider environments in which uncertainty is beneficial and it is not optimal to resolve uncertainty. In these models, the focus is on whether the agent having the information should disclose it to the other agent(s) in the economy. In contrast, the rule designer in our model does not know the policy maker's type and we focus on the design of policies that can induce – or not – revelation. Dovis and Kirpalani (2020a) show that if the rules are chosen by the policy maker, for intermediate levels of discount factors, the commitment type chooses a stringent rule to separate from the optimizing type for all levels of reputation, while the rule designer under the veil of uncertainty chooses to avoid separation when the reputation of the policy maker is sufficiently low.[4]

Our paper is also related to a literature that studies signaling games when policy makers have different types. See, for instance, Vickers (1986), Cole et al. (1995), Angeletos et al. (2006), King et al. (2008), Lu (2013), and Lu et al. (2016) with payoff types, or Dovis and Kirpalani (2020a), where one type has the ability to commit to the announced policy. See also Sanktjohanser (2018) for a similar analysis in the context of a bargaining game. Our approach differs from these papers since we study the best policy chosen by the rule designer when there is uncertainty about the type of the policy maker. As argued above, these two approaches lead to different outcomes.

Debortoli and Nunes (2010) consider a policy game in which the policy maker has the ability to change its policies infrequently and randomly; however, they abstract from reputation-building incentives.

2 Policy game

We consider a policy game that captures a variety of relevant economic environments as special cases. We present two leading examples of our framework: a version of the Barro and Gordon (1983a) model of monetary policy and a banking model in the spirit of Kareken and Wallace (1978). Our framework also nests other models, including the Fisher model of capital income taxation considered in Chari and Kehoe (1990).

There are three types of agents: the *rule designer*, *policy makers* (or bureaucrats), and a continuum of private agents. We consider a repeated environment where there are no physical state variables across periods. At the beginning of each period, the rule designer recommends a policy π_r from a set $[\underline{\pi}, \overline{\pi}]$. We refer to this recommendation as a *rule*. We say that a rule π is more stringent (resp., lenient) than a rule π' if $\pi < \pi'$ (resp.,

[4] In Appendix F.1, we show that the signaling game result in Dovis and Kirpalani (2020a) holds in this model as well.

$\pi > \pi'$). The private agents then choose an individual action. After observing the private action, the policy maker chooses a policy π. The policy maker can be one of two types: a *commitment* type, which always follows the recommendation made by the rule designer, or an *optimizing* type, which can choose any policy π in the set $[\underline{\pi}, \overline{\pi}]$.[5] We assume that the policy maker's type is permanent.[6] The policy maker's type is unobservable to the private agents and the rule designer, who learn about it through the observed policies. We assume that the private agents and the rule designer share a common prior ρ that they are facing the commitment type. We define the probability that the private agents and the rule designer ascribe to the policy maker being the commitment type as the policy maker's *reputation*.

We let x denote the representative (average) action taken by the private agents. We assume that the private action is a function ϕ of the expected policy, $\mathbb{E}\pi = \rho\pi_c + (1-\rho)\,\pi_o$, where $\pi_c = \pi_r$ is the policy chosen by the commitment type and π_o is the policy chosen by the optimizing type,

$$x = \phi\left(\mathbb{E}\pi\right). \tag{1}$$

We will refer to (1) as the *implementability constraint*. We think of the function ϕ as summarizing the set of implementability conditions describing the set of outcomes that can be implemented given a set of policies or an incentive compatibility constraint.

The rule designer and the policy makers maximize a social welfare function $w(x, \pi)$ and discount future payoffs with discount factors β and β_o, respectively. We allow for the rule designer's discount factor β to differ from β_o, although this is not critical. We assume that the problem is time inconsistent. Specifically, we define the Ramsey outcome as

$$(x_{ramsey}, \pi_{ramsey}) = \arg\max_{x,\pi} w(x, \pi) \quad \text{subject to} \quad x = \phi(\pi).$$

We assume that there is a time inconsistency problem in that the Ramsey policy is not optimal ex post, i.e., $\pi_{ramsey} \neq \pi^*(x_{ramsey})$, where $\pi^*(x)$ denotes the best response of the government to x, $\pi^*(x) = \arg\max_\pi w(x, \pi)$. We assume without loss of generality that $\pi^*(x_{ramsey}) > \pi_{ramsey}$.

We also make the following assumptions about w and ϕ:

Assumption 1. *Assume that*

1. *If $w_x > 0$, then $\phi' \leqslant 0$, and $w_{x\pi} < 0$.*

[5] An alternative interpretation is that the type of the policy maker determines the utility cost associated with deviating from the rule designer's recommendation. In particular, we assume that this cost is sufficiently large for the commitment type so that it always follows the recommendation, and zero for the optimizing type.

[6] This assumption is made for convenience. Our main results extend to the case in which the policy maker's type can change exogenously, provided that this type process is persistent. When types are i.i.d., there is no role for reputation.

2. *If* $w_x < 0$, *then* $\phi' \geqslant 0$, *and* $w_{x\pi} > 0$.

As is standard in the time inconsistency literature, we consider environments in which the inability of the policy maker to commit ex post incentivizes the private agents to take worse actions ex ante. Thus, if social welfare is increasing in the private action x, we assume that if the agents expect higher π, they choose lower values of x ($\phi' \leqslant 0$). Finally, we assume a form of supermodularity in (x, π) which implies that the government's incentive to deviate from its ex-ante promises is higher the worse the private action is. In this sense, the private agents' action and the government's static best response are strategic complements.

We next present two economies and show how they map into our general framework.

Example 1: Barro–Gordon One special case of the general environment is the classic Barro and Gordon (1983a) model used to analyze the time-inconsistency problem in monetary policy. In this context, we interpret x as the average wage inflation, and π is the money growth rate (or price inflation).

We assume that the private agents set wage inflation according to

$$x = \phi\left(\mathbb{E}\pi\right) = \rho\pi_c + (1 - \rho)\,\pi_o.$$

The social welfare function takes the quadratic form

$$w\left(x, \pi\right) = -\frac{1}{2}\left[\left(\psi + x - \pi\right)^2 + \pi^2\right],$$

with $\psi > 0$. The first term in this functional form represents the welfare losses associated with low employment, due, for example, to monopolistic competition in labor markets. The parameter ψ measures the extent of this distortion, and it can be mapped into the wage markup set by unions. The second term captures the costs of ex-post inflation, due, for example, to the transactional value of real money balances.

This example captures the problem of designing of the mandate for a central bank. In this context, the rule designer corresponds to congress, while the policy makers are the central bankers who choose the actual policies.

Example 2: Bailout and effort We now consider another economy inspired by the classic analysis in Kareken and Wallace (1978), which studies the trade-off between the ex-post benefits and the ex-ante costs of bailouts. Bailouts are not desirable ex ante because they induce agents to exert inefficiently low effort and in turn reduce expected output; bailouts are ex-post desirable because they avoid bankruptcy costs (net of financing costs).

There are two types of private agents: depositors and bankers. At the beginning of each period, the banker must borrow one unit from the depositors to finance an investment opportunity that pays off at the end of the period. The return on the investment opportunity is R_H with probability $p(e)$, where e is the effort exerted by the banker, and 0 with probability $1 - p(e)$. The function $p(e)$ is increasing and concave, and it satisfies the Inada conditions. Exerting effort e results in a utility cost $\nu(e)$, where ν is increasing and convex. We interpret the banker's effort as the costs associated with monitoring the investment project. The bankers and the depositors are risk-neutral and do not discount consumption between the beginning and the end of the period.

The banker offers the depositors a contract that promises to repay R units of the consumption good in the second sub-period subject to limited liability. We assume that society faces bankruptcy costs ψ whenever the lenders recover less than their initial investment.[7] The policy maker can avoid these bankruptcy costs by making a transfer to the banker so that the depositors can be repaid. In particular, the government can choose the recovery π in case the banker is unable to repay. There is a taxation cost associated with these transfers, denoted by $c(\pi)$, where c is increasing and convex.

We assume that the depositors can observe the effort e. The interest rate schedule, $R(e)$, faced by a banker is determined by the depositors' break-even condition, which requires that the expected return from lending equal 1,

$$p(e) R(e) + (1 - p(e)) [\rho\pi_c + (1-\rho)\pi_o] = 1. \tag{2}$$

Note that for a given effort level, higher expected bailouts in case of default, $\mathbb{E}\pi = \rho\pi_c + (1-\rho)\pi_o$, result in a lower interest rate because depositors require a lower payment in the no-default state given the larger expected payment in the default state. The banker chooses the effort to maximize $-\nu(e) + p(e)[R_H - R(e)]$, where the interest rate schedule, $R(e)$, is implicitly defined by (2). Using (2) to substitute for $R(e)$, we can rewrite the banker's problem as

$$\max_e -\nu(e) + p(e) R_H + (1 - p(e)) \mathbb{E}\pi,$$

where the term $(1 - p(e))\mathbb{E}\pi$ represents the distortion induced by the expected bailout. Thus the optimal effort e is a function $\phi(\mathbb{E}\pi)$ that is implicitly defined by the banker's first-order condition

$$\nu'(e) = p'(e)[R_H - \mathbb{E}\pi].$$

The social welfare function is the equally weighted sum of the utility of the bankers and

[7] Alternatively, we could have assumed that these costs are incurred whenever the lenders recover less than the promised return.

depositors net of taxation ($c(\pi)$) and bankruptcy costs ($\psi(1-\pi)$):

$$w(e,\pi) = -\nu(e) + p(e) R_H - 1 - (1-p(e))(1-\pi)\psi - c(\pi).$$

To simplify calculations we assume that $p(e) = e^{\alpha}$, $\nu(e) = e^2/2$, and $c(\pi) = \lambda\pi^2/2$.

This example captures the problem of designing optimal financial regulation to be implemented by a regulatory authority. A recent example is the Single Resolution Mechanism in the European Union. The Parliament and the Council of the European Union approved the regulation that is then implemented by the Single Resolution Board (SRB). Thus, the Parliament and the Council correspond to the rule designer, and the SRB to the policy maker in our model.

3 Statically optimal rules

We now consider the problem of how to design the optimal rule in a static environment. We establish a set of sufficient conditions under which uncertainty about the policy maker's type is beneficial in the static economy.

Setup The rule designer anticipates that if the policy maker is the commitment type, it will follow the rule π_r. Instead, if the policy maker is the optimizing type, it will always choose the static best response to the private action x. This is because in a static model the rule designer has no tools to incentivize the optimizing type to take any other action. Of course, this will change in the dynamic setting.

To characterize the equilibrium, we write the rule designer's problem as if it directly chooses the equilibrium outcomes subject to the appropriate incentive compatibility constraints. In particular, the rule designer solves

$$W_0(\rho) = \max_{\pi_c,\pi_o,x} \rho w(x,\pi_c) + (1-\rho) w(x,\pi_o), \tag{3}$$

subject to the implementability constraint,

$$x = \phi(\rho\pi_c + (1-\rho)\pi_o), \tag{4}$$

and the optimizing type's incentive constraint,

$$\pi_o = \pi^*(x).$$

For later reference, we denote the solution to this problem as $\pi_{c,0}(\rho)$, $\pi_{o,0}(\rho) = \pi^*(x_0(\rho))$, and $x_0(\rho)$.[8] We also define the value for the optimizing type:

$$V_0(\rho) = w(x_0(\rho), \pi_{o,0}(\rho)).$$

Uncertainty is beneficial We next discuss the conditions under which *uncertainty is beneficial* in that

$$W_0(\rho) \geqslant \rho W_0(1) + (1-\rho) W_0(0). \tag{5}$$

When uncertainty is beneficial, the expected social welfare is higher when the policy maker's type is uncertain relative to the case in which types are revealed right before the rule designer chooses the rule. This property of the static problem turns out to be critical for the form of the optimal rule in a dynamic model.

We now provide a set of sufficient conditions on primitives that ensure that uncertainty is beneficial. In the appendix we show that the Barro–Gordon and bailout examples satisfy these assumptions.

Assumption 2. *Assume:*

1. *Concave preferences:* $w(x,\pi)$ *is concave in* (x,π).
2. *Strategic complementarities decreasing in reputation:* $w_\pi(x,\pi)$ *is convex in* (x,π) *and if* $w_x > 0$ *then* $\phi'' \leqslant 0$ *or if* $w_x \leqslant 0$ *then* $\phi'' \geqslant 0$.
3. $w_\pi(x,\pi) + [\rho w_x(x,\pi) + (1-\rho) w_x(x,\pi^*(x))] \frac{\phi'(\cdot)}{[1-\phi'(\cdot)(1-\rho)\pi^*_x(x)]} \leqslant 0$ *for all* π.
4. $1 > \pi^*{}_x(x)\phi'(\pi) \geqslant 1 - \frac{w_x(\underline{x},\pi^*(\underline{x}))}{w_x(\underline{x},\underline{\pi})}$, *where* $x = \phi(\pi)$, $\underline{x} = \phi(\underline{\pi})$, *and* $\pi^*{}_x(x) = -\frac{w_{x\pi}(x,\pi^*(x))}{w_{\pi\pi}(x,\pi^*(x))}$.

We have the following result:

Proposition 1. *Under Assumptions 1 and 2,* $\pi_{c,0}(\rho) = \underline{\pi}$ *and uncertainty is beneficial in that (5) holds.*

The proof of this, and of all other propositions, is provided in the appendix. Condition 3 in Assumption 2 implies that the optimal static rule takes a simple form: for all ρ, the rule is set to $\pi_c(\rho) = \underline{\pi}$, which is also the Ramsey policy.[9] This is true even though the private action is not at the Ramsey level since the private agents anticipate that with probability $1-\rho$ the policy maker is the optimizing type, who will deviate from

[8] Note that here we are allowing the rule designer to choose the best equilibrium given a rule π_r. Thus there is no need to have the rule depend on the representative private action x.

[9] Clearly, for $\rho = 0$ the optimal rule is a correspondence with values $[\underline{\pi}, \bar{\pi}]$ because the right side of (3) is independent of π_c.

the recommendation and choose the static best response. This is because the expression in Condition 3 is the first-order condition for the problem in (3). The condition implies that reducing π has a positive marginal effect and thus it is optimal to be at the corner $\underline{\pi}$. Therefore, it is optimal for the rule designer to recommend the most stringent policy. Intuitively, the rule designer is trading off the cost of not best responding to the private action (the first term in Condition 3) with the benefit of inducing private agents to take a more favorable action (the second term in Condition 3). Under Condition 3, the benefit outweighs the cost for all levels of the private action, even those away from the Ramsey outcome. Thus, it is optimal to choose the most stringent policy. In the context of the Barro–Gordon example, this means that the optimal inflation target is zero, while in the bailout example, a strict no-bailout policy is optimal. Restricting to environments where Condition 3 holds simplifies the analysis without changing the economics.

The critical role of the Proposition is to identify the sufficient conditions such that condition (5) holds. Here we sketch the logic of our proof for the case in which $w_x > 0$, as in the bailout model. A specular logic holds for the case in which $w_x < 0$. To establish that uncertainty is beneficial it is sufficient to show that $W_0(\rho)$ is concave, or, equivalently, that there are decreasing returns to reputation. First, we prove that $w(x_0(\rho), \underline{\pi})$ and $w(x_0(\rho), \pi^*(x_0(\rho)))$ are concave in ρ. To this end, note that since ϕ is concave (Condition 2 of Assumption 2) and π^* is decreasing and convex (which follows from Condition 2 of Assumption 2), then $x_0(\rho) = \phi(\rho\underline{\pi} + (1-\rho)\pi^*(x_0(\rho)))$ is concave in ρ. These observations together with the concavity of w and $w_x > 0$ imply that $w(x_0(\rho), \underline{\pi})$ and $w(x_0(\rho), \pi^*(x_0(\rho)))$ are concave in ρ. Having established the concavity of $w(x_0(\rho), \underline{\pi})$ and $w(x_0(\rho), \pi^*(x_0(\rho)))$ is not enough to show that

$$W_0(\rho) = \rho w(x_0(\rho), \underline{\pi}) + (1-\rho) w(x_0(\rho), \pi^*(x_0(\rho)))$$

is concave since the product of two concave functions is not necessarily concave. However, the technical assumption in Condition 4 guarantees that $W_0(\rho)$ is concave.

There are two key economic conditions that deliver the concavity of $W_0(\rho)$. First, the rule designer's preferences are concave; this is Condition 1 in Assumption 2. This condition implies that the costs of deviating from the Ramsey outcome are convex. Second, the private action x and the policy π_o are strategic complements – in that ϕ is decreasing in π_o and π^* is decreasing in x – and this complementarity is stronger when the private action x is far from the Ramsey outcome (low level of reputation); this is Condition 2 in Assumption 2. This condition is illustrated in Figure 1. The first panel shows that the change in the private action to a decrease in π_o is larger when the expected policy is high (far from Ramsey policy), i.e., ϕ is concave. The second panel shows that the change in the policy maker's best response to an increase in the private action is large (in absolute

Figure 1: Static value and private action when $w_x > 0$

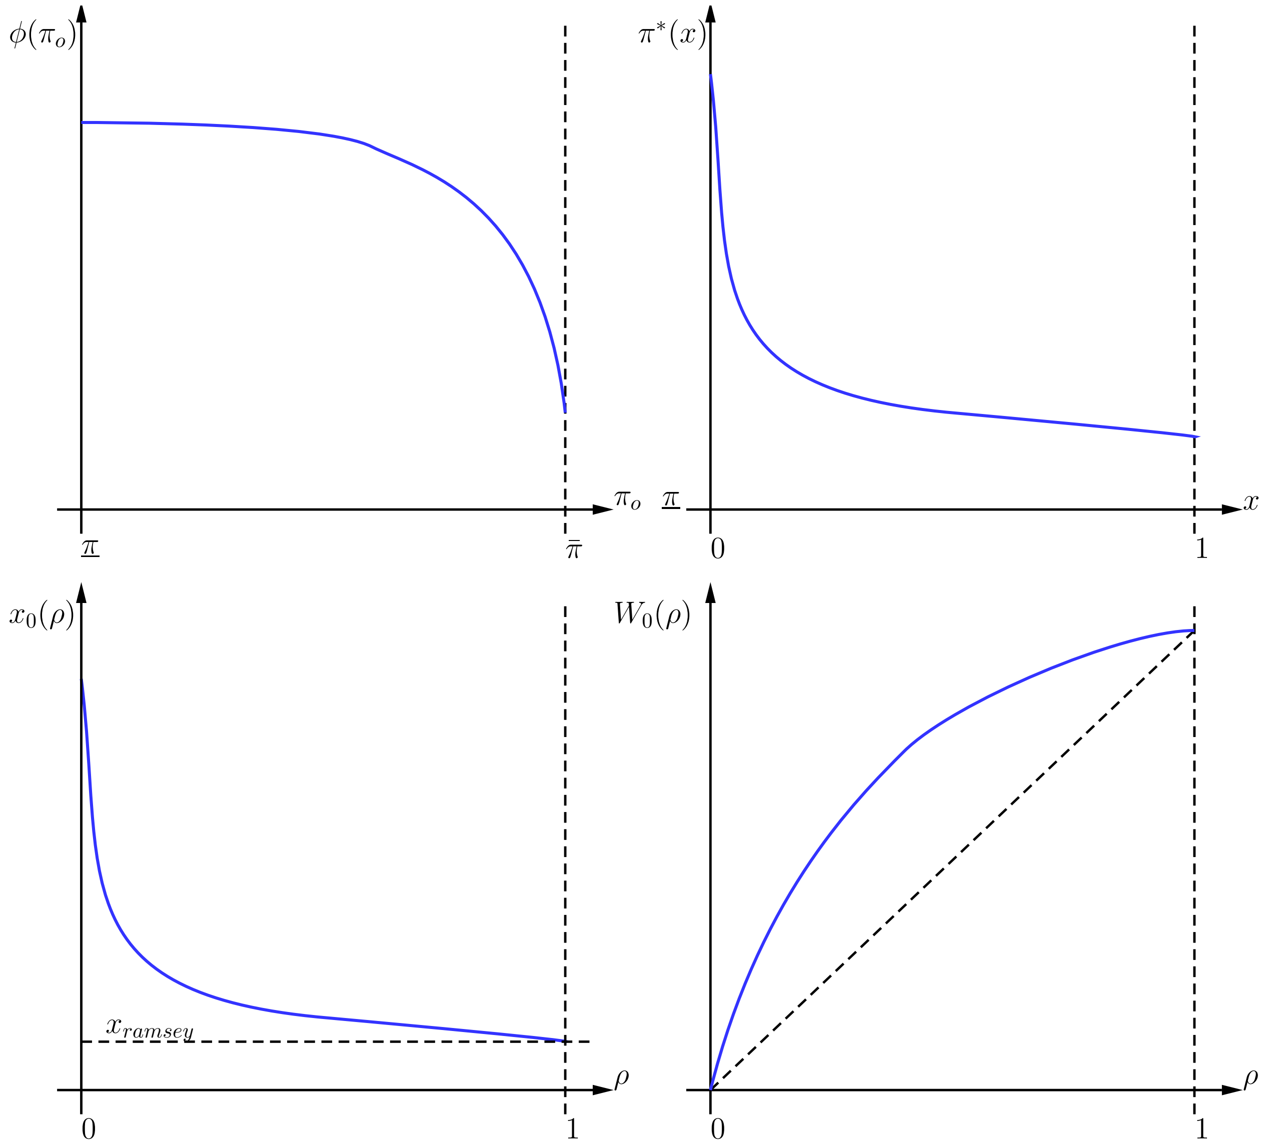

value) when the action is low (far from the Ramsey outcome), i.e., π^* is convex. This implies that an increase in reputation has a larger effect on the average private action when reputation is low (x_0 is concave), as shown in the third panel of Figure 1.

These two conditions are critical to ensuring that there are decreasing returns to reputation. Note, however, that the two conditions are neither sufficient nor necessary, although at least one of them must hold.

Discussion of key assumptions We now discuss the properties of the underlying economic environments in which the aforementioned conditions – concave preferences and stronger complementarities when reputation is low – are likely to be met. Clearly, the majority of applications assume concave preferences. For instance, in monetary models, the cost of inflation is assumed to be convex, as in the Barro–Gordon example.

Consider now the strategic complementarities condition. For the best response of the optimizing type to be (weakly) convex, it must be that w_π is convex. Intuitively, the time inconsistency problem is exacerbated when the private action x is farther away from the Ramsey outcome.

Next, consider the concavity of ϕ. So far, we have taken ϕ to be a primitive object in our policy game. To discuss the underlying properties that give a concave ϕ, note that in general we can think of ϕ as arising from the following maximization problem by an individual private agent i:

$$\max_{x_i} u\left(x_i, x, \pi\right)$$

that takes as given the average action x and the policy π (which is a function of the average action x) for some indirect utility function u. Assuming concavity of u in x_i, optimality requires that $u_1\left(x_i, x, \pi\right) = 0$. By imposing representativeness we can then derive $\phi\left(\pi\right)$ as the solution to

$$u_1\left(\phi\left(\pi\right), \phi\left(\pi\right), \pi\right) = 0. \tag{6}$$

Recall that this was exactly how we derived ϕ for the bailout example.

For ϕ to be concave, it must be that the incentives to choose an action far away from the Ramsey action are higher if private agents expect the policy to be far away from the Ramsey policy. If the average action does not directly affect the agent's utility in that $u\left(x_i, x, \pi\right) = u\left(x_i, \pi\right)$, as is the case in our two examples, this is true if u_1 is jointly concave in $\left(x_i, \pi\right)$. When the indirect utility u also depends on x – say because of the presence of a price that in equilibrium depends on x – then one must also consider the complementarity between x_i and x. For instance, if x_i and x are substitutable – $u_{13} < 0$ – and this effect is stronger the higher is x, then ϕ may be convex in ρ and it does not satisfy our conditions.

In Appendix B, we provide an example where this is the case based on a simple investment model in which the government can expropriate the investment made by foreign investors ex post but ex ante has incentives to promise no taxes to stimulate investment. In the example, the foreign investment, $x\left(\pi\right)$, is increasing and convex in expected taxes π. When (expected) taxes are high, a reduction in taxes results in a smaller increase in investment than when taxes are low. That is, there are increasing returns to reputation. This is because the equilibrium marginal product of capital is more sensitive when taxes are high and thus the effect of the tax reduction on investment is mitigated by the steep reduction in the marginal product of capital. Thus, the complementarity between taxes and investment is lower when taxes are high (far away from the Ramsey policy). This in turn implies that $W_0\left(\rho\right)$ is convex if the government's preferences for consumption are close to linear.

4 Dynamically optimal rules

We next study how the optimal rule changes once we introduce dynamics. Our main result is that when reputation is low, the rule designer chooses a rule which preserves uncertainty about the type of the policy maker. The optimal rule is the most stringent policy that is incentive compatible for the optimizing type. This recommended policy is more lenient than the statically optimal one. Leniency in the rule makes it easier for the optimizing type to follow the recommendation ex post. In contrast, if reputation is sufficiently high, the rule designer finds it optimal to set stringent rules that result in the type of the policy maker being revealed. This is because when reputation is sufficiently high, the static costs associated with choosing a lenient rule outweigh the benefits associated with preserving uncertainty about the policy maker's type.

We assume that the rule designer can choose the optimal rule in each period without commitment. In Appendix E, we show that in the twice repeated economy, the solutions with and without commitment on the part of the rule designer coincide. With more than two periods, whether these two values coincide depends on the level of reputation. We show that for reputation values that are either sufficiently high or sufficiently low, the commitment and no-commitment outcomes coincide. Thus, for reputation values in these ranges, our main results are unchanged if the rule designer has commitment.[10]

4.1 Two-period problem

We now study the optimal rule design problem in a two-period economy by repeating the stage game studied in the previous section twice. When there is more than one period, the optimizing type can be incentivized to take a different action from its static best response. We can set up the rule designer's problem as choosing the rule that will be followed by the commitment type, $\pi_r = \pi_c$, and a recommendation to the optimizing type, π_o. This recommendation must be incentive compatible in that the optimizing type must prefer to follow the recommendation than to choose its best possible deviation (playing the static best response $\pi^*(x)$) and attaining a continuation value $V_0(0)$ as the prior jumps to zero:

$$w(x, \pi_o) + \beta_o V_0 \left(\rho' (\pi_o | \pi_c, \sigma) \right) \geqslant w(x, \pi^*(x)) + \beta_o V_0(0), \tag{7}$$

where recall β_o is the discount factor for the optimizing type and $\rho'(\pi_o|\pi_c, \sigma)$ is the private agents' belief about the policy maker's type after observing π_o given recommenda-

[10]However, for an intermediate range of reputation values, the solution to the problem with commitment differs from the solution in the case in which rules are chosen sequentially: the rule designer itself suffers from a time inconsistency problem. This is because future rules can be used to incentivize the policy maker in the current period, thereby relaxing the incentive compatibility constraint.

tion π_c and the indicator variable σ, which takes value 1 if the private agents expect the optimizing type to choose the same policy as the commitment type, $\pi_o = \pi_c$, and $\sigma = 0$ otherwise.[11] This function for the law of motion for beliefs on path follows Bayes' rule and is given by

$$\rho'(\pi|\pi_c,\sigma) = \begin{cases} \frac{\rho}{\rho+(1-\rho)\sigma} & \text{if } \pi = \pi_c \\ 0 & \text{o/w} \end{cases}. \tag{8}$$

For all subsequent analyses we assume that the policy makers are sufficiently impatient so that the Ramsey outcome is not incentive compatible:

Assumption 3. *The discount factor β_o is small enough so that*

$$w(x_{ramsey}, \pi^*(x_{ramsey})) - w(x_{ramsey}, \pi_{ramsey}) > \frac{\beta_o}{1-\beta_o}\left[V_0(1) - V_0(0)\right].$$

As in the static case, we characterize the equilibrium by solving the problem of the rule designer that chooses the equilibrium outcome subject to the appropriate incentive constraints. The rule designer's problem is[12]

$$\begin{aligned} W(\rho) = \max_{x,\pi_c,\pi_o,\sigma} \rho & \left[w(x,\pi_c) + \beta W_0\left(\rho'(\pi_c|\pi_c,\sigma)\right)\right] \\ & + (1-\rho)\left[w(x,\pi_o) + \beta W_0\left(\rho'(\pi_o|\pi_c,\sigma)\right)\right] \end{aligned} \tag{9}$$

subject to the implementability condition, (4), and the incentive compatibility constraint for the optimizing type, (7), and $\pi_o = \sigma\pi_c + (1-\sigma)\pi^*(x)$, given the law of motion of beliefs $\rho(\pi|\pi_c,\sigma)$ defined in (8).

For simplicity, we abstract from mixed strategies for the optimizing type. In Appendix C.2, we show that this is without loss of generality in this two-period economy. Under our assumptions, the outcome in which the optimizing type follows the rule with probability $\sigma \in (0,1)$ and the ex-post optimal policy with probability $1-\sigma$ is dominated in terms of welfare by the best equilibrium in which the optimizing type follows the rule with probability one. This is because of two reasons. First, since uncertainty is beneficial and the posterior is a martingale, mixing introduces volatility in the posterior without affecting its mean, which lowers the rule designer's expected continuation value. Second, as we show in the appendix, mixing tightens the optimizing type's incentive constraint because w is concave in π and V_0 is concave in ρ and thus reduces static payoffs.

[11] If the rule designer chooses $\pi_o = \pi_c$, then a deviation only happens off-path and thus Bayes' rule does not pin down the posterior. On the right side of (7), we assume that after a deviation, the posterior goes to zero. This is reasonable because the commitment type cannot deviate. Moreover, it also constitutes the worst punishment if the optimizing type deviates.

[12] In Appendix C, we define an equilibrium for the policy game and show that the problem in (9) characterizes the best Perfect Bayesian Equilibrium outcome.

We can then reduce the problem above to a discrete choice between two options: separating or pooling. If the rule designer chooses to separate, it chooses the best static rule. Because of Assumption 3, the Ramsey outcome is not incentive compatible and the optimizing type will choose the static best response and not follow the rule so the type of the policy maker is revealed at the end of the period. Thus the continuation value is either $W_0(1)$ with probability ρ or $W_0(0)$ with probability $1-\rho$. The value of separating is then

$$W_{sep}(\rho) = W_0(\rho) + \beta\left[\rho W_0(1) + (1-\rho) W_0(0)\right].$$

If the rule designer chooses to pool, it sets the rule to $\pi_{ico,1}(\rho)$, which is the most stringent policy π consistent with the incentive compatibility constraint for the optimizing type:

$$w(\phi(\pi_{ico,1}(\rho)), \pi_{ico,1}(\rho)) + \beta V_0(\rho) = w(\phi(\pi_{ico,1}(\rho)), \pi^*(\phi(\pi_{ico,1}(\rho)))) + \beta V_0(0).$$

In this case, both types of policy makers follow the rule in equilibrium and thus uncertainty about their type is preserved and the continuation value is $W_0(\rho)$. The value of pooling is then

$$W_{pool}(\rho) = w(\phi(\pi_{ico,1}(\rho)), \pi_{ico,1}(\rho)) + \beta W_0(\rho).$$

The next proposition shows that designing a rule that preserves uncertainty about the policy maker's type is valuable when its reputation is low:

Proposition 2. *Under Assumptions 1–3*, there exist ρ_1^* and ρ_2^* with $0 < \rho_1^* \leqslant \rho_2^* < 1$ such that:

1. *For* $\rho \in \left[\rho_2^*, 1\right]$ *there is separation* ($\sigma = 0$) *and* $\pi = \underline{\pi}$.
2. *For* $\rho \in [0, \rho_1^*]$ *there is pooling* ($\sigma = 1$) *and* $\pi_c(\rho) > \underline{\pi}$.

The key implication of this proposition is that in contrast to the static case, when reputation is low, the rule designer recommends more lenient rules in order to preserve uncertainty about the policy maker's type in the future. To see why this is the case, consider

$$W_{pool}(\rho) - W_{sep}(\rho) = \Delta\omega(\rho) + \beta\Delta\Omega(\rho),$$

where $\Delta\Omega(\rho) \equiv W_0(\rho) - [\rho W_0(1) + (1-\rho) W_0(0)]$ are the *dynamic benefits of pooling* and $\Delta\omega(\rho)$ are the *static benefits of pooling* given by

$$\Delta\omega(\rho) \equiv w(\phi(\pi_{ico,1}(\rho)), \pi_{ico,1}(\rho)) - W_0(\rho).$$

The dynamic and static benefits of pooling are plotted in Figure 2. Since uncertainty is beneficial, we know that $\Delta\Omega(\rho) > 0$ for all $\rho \in (0,1)$ and equal to zero when there is no uncertainty and $\rho \in \{0,1\}$. Also, by construction, the static benefits of pooling are zero for $\rho = 0$, since $\pi_{ico,1}(0) = \pi_{o,0}(0) = \pi^*(x_0(0))$, and negative for $\rho = 1$, since $W_0(1)$ attains the Ramsey value and $w(\phi(\pi_{ico,1}(\rho)), \pi_{ico,1}(\rho)) < W_{ramsey}$ because the incentive constraint is assumed to be binding for all ρ (Assumption 3).

Figure 2: **Dynamic and static benefits of pooling**

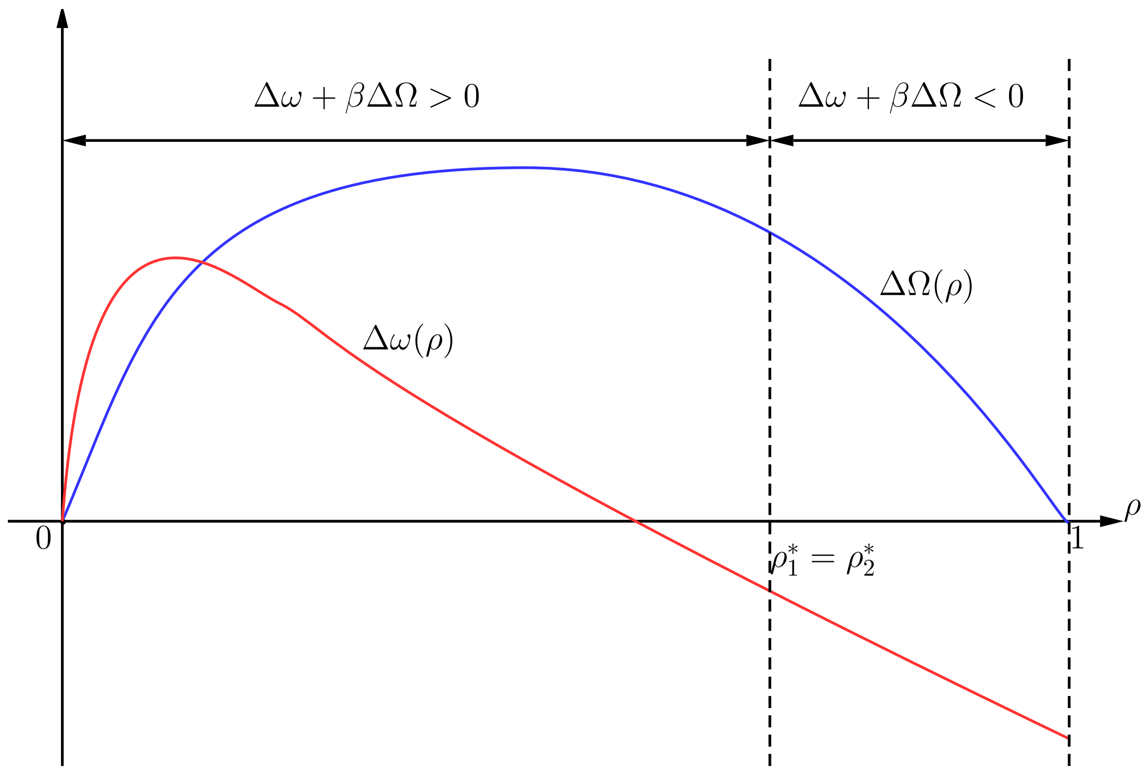

Combining these observations, it is immediate that for ρ close to one, $W_{sep}(\rho) > W_{pool}(\rho)$ since the dynamic benefits are approximately zero and $\Delta\omega(\rho) < 0$. In the proof, we show that the static benefits of pooling $\Delta\omega(\rho)$ are increasing in ρ for low levels of reputation. Intuitively, in the pooling regime the rule designer induces the optimizing type to follow a more stringent policy than the static best response, $\pi_{ico,1}(\rho) < \pi^*(x_0(\rho))$ at the cost of forcing the commitment type to follow a more lenient policy, $\pi_{ico,1}(\rho) > \underline{\pi}$. When reputation is low enough, the expected policy becomes more stringent in the pooling regime than in the separating regime, because in the latter, the private agents expect the recommended policy to be followed with a low probability. Thus, pooling has both static and dynamic benefits and is therefore preferable to separation.[13]

[13]Note that the positive static benefits of pooling for low levels of reputation do not depend on uncertainty being beneficial. Thus pooling may be optimal for low levels of reputation even in economies that do not satisfy our assumptions and thus have a convex continuation value $W_0(\rho)$ and negative dynamic benefits of pooling.

For the Barro–Gordon example we can provide a tighter characterization of the optimal policy and show that it has a cutoff property, i.e., $\rho_1^* = \rho_2^*$. The proof for this is in the appendix. (For the bailout example we verify that this is the case numerically.) The optimal dynamic rule in this case is plotted in Figure 3.

Figure 3: **Optimal dynamic rule**

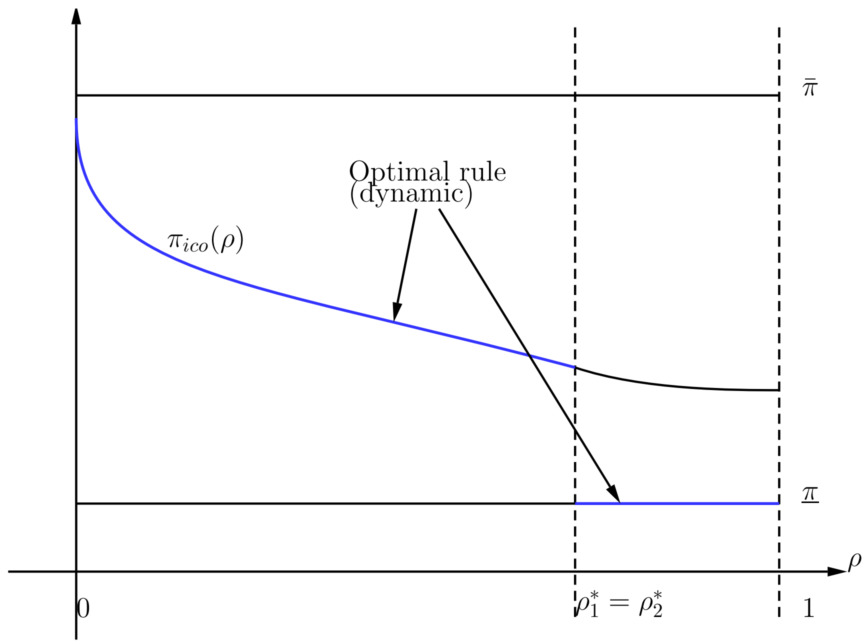

Let us now consider what Proposition 2 implies for our two examples. In the bailout example, the optimal static rule is a strict no-bailout policy. However, in the dynamic model, on-path bailouts are necessary to achieve good outcomes when reputation is low. In particular, counter to conventional wisdom, bailouts along the equilibrium path are necessary in order to impose future discipline on financial institutions. This is precisely because allowing for bailouts makes it easier for the optimizing type to follow the designer's recommendation and thus helps to preserve uncertainty going forward. This is beneficial because uncertainty about the policy maker's type prevents bankers from taking on excessive risk by exerting little effort. Similarly, in the Barro–Gordon model, having looser inflation targets is beneficial when reputation is low.

4.2 Limit of finite horizon

We now show that the insights from the two-period model extend to any horizon. In particular, we analyze the limit of the finite-horizon economy and show that an analog of Proposition 2 holds.

Let k denote the horizon of the economy, i.e., the number of periods left. Let $\{\pi_k(\rho)\}_{k=0}^{\infty}$ be the optimal rules set by the rule designer at each horizon as a function of the prior ρ.

In the previous sections, we characterized the case for $k = 0, 1$. We will use the following property:

Assumption 4. *The gains of best responding are decreasing in* x*, in that*

$$G(x) \equiv w(x, \pi^*(x)) - w\left(x, \phi^{-1}(x)\right)$$

is monotone decreasing in x*.*

This property is satisfied in our two examples. In the appendix we provide an additional sufficient condition on the general environment which implies this property.

To set up our next proposition, we define the following objects. First, let (x_{CK}, π_{CK}) be the private action and the policy that emerge in the best sustainable equilibrium for the infinite-horizon version of the model where $\rho = 0$. That is, x_{CK} solves

$$\frac{w\left(x_{CK}, \phi^{-1}(x_{CK})\right)}{1-\beta_o} = w(x_{CK}, \pi^*(x_{CK})) + \frac{\beta_o}{1-\beta_o} W_0(0), \tag{10}$$

where $W_0(0) = V_0(0)$ is the value of the worst equilibrium (the repetition of the static Nash for $\rho = 0$) and $\pi_{CK} = \phi^{-1}(x_{CK})$. Note that because of Assumption 3, x_{CK} is higher (resp., lower) than the Ramsey outcome when $w_x < 0$ (resp., $w_x > 0$).

The next proposition shows that the limit of the finite-horizon economy has the following property: there are two cutoffs and it is optimal to pool for priors below one cutoff and separate for priors above the other cutoff.

Proposition 3. *Under Assumptions 1–4, as the horizon* $k \to \infty$ *we have that:*

1. *For* $\rho = 0$*,* $W_k(0) \to W_0(0)/(1-\beta)$ *and* $V_k(0) \to V_0(0)/(1-\beta_o)$*.*
2. *There exists* $\hat{\rho} \in (0, \rho_1^*)$ *such that for* $\rho \in (0, \hat{\rho}]$*, there is pooling for all* k *and* $\pi_k(\rho) \to \pi_{CK}$*.*
3. *For* $\rho \in (\rho_2^*, 1]$*, there is separation for all* k *and* $\pi_k(\rho) = \underline{\pi}$ *for all* k*.*

Qualitatively, the optimal rule is the same as in the two-period model. For high values of ρ, above the cutoff ρ_2^*, it is optimal to separate because pooling is associated with static losses that are not compensated by the dynamic gains. For low levels of reputation, it is optimal to choose rules that do not reveal the type of the policy maker. Note that the optimal policy in the pooling regime does not depend on the prior ρ in the limit. This is because if it is optimal to pool today it is also optimal to pool in all subsequent periods. In this case, the type of the policy maker will never be revealed and so ρ does not affect the value on the equilibrium path. The initial prior also does not affect the value of the deviation because the posterior jumps to zero independently of the initial value. Thus the value of pooling in $(0, \hat{\rho}]$ is independent of ρ, as shown in Figure 4. Moreover, the policy converges to its value in the best sustainable equilibrium when it is known that the policy maker is the optimizing type.

Notice that there is a discontinuity at $\rho = 0$. This is because when $\rho = 0$ and the horizon is finite it is not possible to incentivize the optimizing type to choose any policy other than its static best response.

For intermediate values, $\rho \in (\hat{\rho}, \rho_1^*)$, the equilibrium strategies may not converge as the horizon goes to infinity, and randomization may be optimal.

Figure 4: **Equilibrium values**

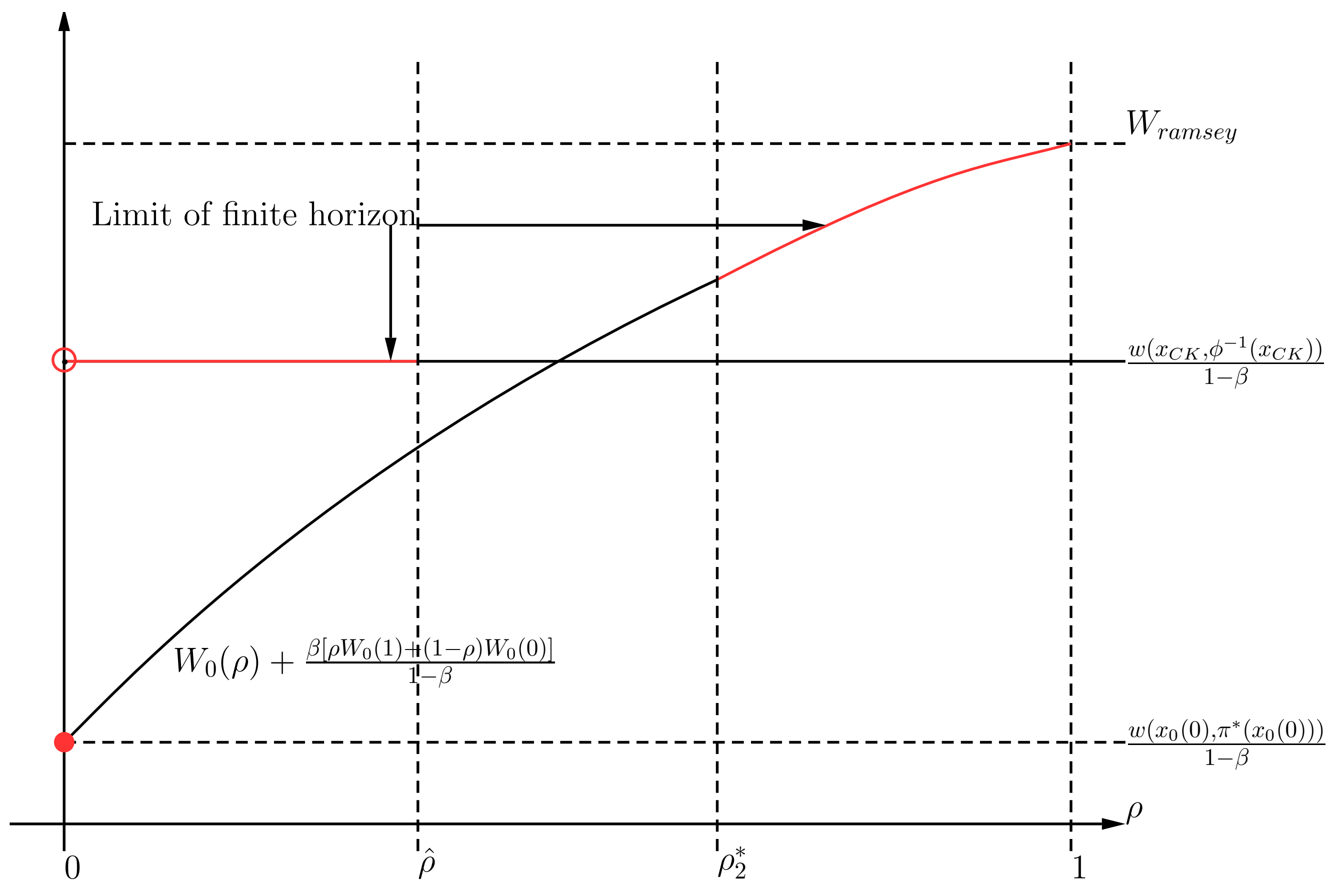

So far we have allowed the rule designer to choose a new rule in each period as a function of the current reputation of the policy maker. In practice, opportunities for revising and introducing new rules arise infrequently. Suppose for instance that rules are "sticky" in that they can only be changed in a given period with probability $\alpha < 1$. The inability to change rules for sure next period affects the rule designer's value when it chooses to pool. In particular, choosing π_c equal to the lowest value consistent with the optimizing type's incentive compatibility constraint induces the optimizing type to pool in the current period, but it may induce separation in future periods if the rules cannot be adjusted. Thus, the rule designer may want to choose an even more lenient policy, which implies that the optimizing type's incentive constraint is slack, to ensure pooling next period in the event that the rule cannot be adjusted. However, this trade-off vanishes in the limit for $\rho \in [0, \hat{\rho}]$ as the horizon goes to infinity since $\lim_{K\to\infty} \pi^K_{ico,t}(\rho) = \lim_{K\to\infty} \pi^K_{ico,t+1}(\rho) = \pi_{CK}$. Thus the minimal rule that ensures pooling is constant over time. This observation implies that with sticky policies ($\alpha < 1$) there exists an equilibrium with the same properties outlined in Proposition 3.

5 Transparency of rules

We now study the implications of our theory for the optimal degree of transparency of the rule. Should the rule be designed so that a deviation by the policy maker is easily detectable? In other words, we ask if perfect monitoring is always desirable. In repeated policy games with no reputational considerations, perfect monitoring is always desirable; see Atkeson and Kehoe (2001), Atkeson et al. (2007), and Piguillem and Schneider (2013). In contrast, we show that with reputational considerations, perfect monitoring is desirable only for low levels of reputation, while imperfect monitoring is desirable for high levels of reputation.

5.1 Optimal degree of monitoring

We first consider the case in which the rule designer can control the degree to which the private agents and future rule designers can monitor the policies chosen by the policy maker. In particular, suppose the private agents cannot directly observe the policy π, but they can only observe a signal $\tilde{\pi} = \pi + \varepsilon$, where $\varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$. The rule designer can choose the standard deviation of the noise, σ_ε, as part of the optimal rule design. We interpret the choice of large noise as standing in for complicated rules whose deviations are hard to detect for the private agents. We say that a rule is *transparent* if $\sigma_\varepsilon = 0$ and *opaque* if $\sigma_\varepsilon > 0$.

For a given σ_ε, the law of motion for beliefs is

$$\rho'\left(\tilde{\pi}|\rho, \sigma_\varepsilon, \pi_c, \pi_o\right) = \frac{\rho g\left(\tilde{\pi} - \pi_c|\sigma_\varepsilon\right)}{\rho g\left(\tilde{\pi} - \pi_c|\sigma_\varepsilon\right) + (1-\rho)\, g\left(\tilde{\pi} - \pi_o|\sigma_\varepsilon\right)}, \tag{11}$$

where g is the PDF of a Normal distribution with mean zero and variance σ_ε^2. We can then write the rule designer's problem for the twice repeated economy as

$$\max_{x,\pi_c,\pi_o,\sigma_\varepsilon} \rho \left[w(x,\pi_c) + \beta \int W_0\left(\rho'(\pi_c+\varepsilon)\right) g(\varepsilon)\, d\varepsilon \right] + (1-\rho)\left[w(x,\pi_o) + \beta \int W_0\left(\rho'(\pi_o+\varepsilon)\right) g(\varepsilon)\, d\varepsilon \right] \tag{12}$$

subject to the implementability constraint, $x = \rho\pi_c + (1-\rho)\pi_o$, and the incentive compatibility constraint for the optimizing type,

$$w(x,\pi_o) + \beta_o \int V_0\left(\rho'(\pi_o+\varepsilon,\rho)\right) g(\varepsilon)\, d\varepsilon \geqslant w(x,\pi) + \beta_o \int V_0\left(\rho'(\pi+\varepsilon,\rho)\right) g(\varepsilon)\, d\varepsilon \quad \forall \pi, \tag{13}$$

taking as given the law of motion for beliefs defined in (11). Note that the values in the final period, W_0 and V_0, are the static values and are not affected by σ_ε.

The next proposition establishes that for low levels of reputation it is optimal to have perfectly transparent rules ($\sigma_\varepsilon = 0$), while for higher values of reputation it is optimal to have opaque rules. Let ρ_1^* and ρ_2^* be the cutoffs defined in Proposition 2:

Proposition 4. *Under Assumptions 1–4:*

1. *For $\rho \in \left[0, \rho_1^*\right]$ there is pooling and signals are perfectly informative, $\sigma_\varepsilon = 0$.*
2. *For $\rho \in \left[\rho_2^*, 1\right]$ there is separation and signals are not perfectly informative, $\sigma_\varepsilon > 0$.*

Consider first low levels of reputation. From Proposition 2, we know that if signals are perfectly informative, it is optimal to be in the pooling regime, so $\pi_o = \pi_c$. Conditional on pooling, it is preferable to choose $\sigma_\varepsilon = 0$ to relax the incentive constraint (13). In fact, without noise, (13) reduces to

$$w(x,\pi_o) + \beta_o V_0(\rho) \geqslant w(x,\pi) + \beta_o V_0(0) \quad \forall \pi, \tag{14}$$

and so the spread in continuation values $[V_0(\rho) - V_0(0)]$ provides the maximal incentives to the optimizing type. To see this, first note that for any $\sigma_\varepsilon > 0$

$$\int V_0\left(\rho'(\pi+\varepsilon,\rho)\right) g(\varepsilon)\, d\varepsilon > V_0(0),$$

so the right side of (13) is lowest at $\sigma_\varepsilon = 0$. Second, by concavity of V_0 we have that

$$V_0(\rho) > \int V_0\left(\rho'(\pi_o+\varepsilon,\rho)\right) g(\varepsilon)\, d\varepsilon$$

since $\rho = \int \rho'(\pi_o+\varepsilon,\rho)\, g(\varepsilon)\, d\varepsilon$, so the left side of (13) is highest at $\sigma_\varepsilon = 0$. Thus, since we know that for low levels of reputation pooling is preferable to separating, we have that the optimal rule has pooling and it is perfectly transparent.

Consider now high levels of reputation. From Proposition 2, we know that for $\rho \in [\rho_2^*, 1]$ separation is preferable to pooling when signals are perfectly informative. As argued above, conditional on pooling, it is preferable to choose $\sigma_\varepsilon = 0$ to relax the incentive constraint (13). Thus, for these high levels of reputation, pooling cannot be optimal and we are in the separating regime ($\pi_o \neq \pi_c$). Suppose by way of contradiction that it is optimal to have perfectly informative signals, $\sigma_\varepsilon = 0$. Since types are perfectly revealed at the end of the first period, we have that $\rho' \in \{0, 1\}$ and the only incentive-compatible policy for the optimizing type is $\pi_o = \pi^*(x)$. Note that we can support the same policies by choosing $\sigma_\varepsilon = \infty$. This alternative rule has the same static payoff but prevents learning about the regulator's type and therefore $\rho' = \rho$, because the signal $\tilde{\pi}$ is totally uninformative. This increases the expected continuation value because uncertainty is beneficial, $W(\rho) > \rho W(1) + (1-\rho) W(0)$. Thus, the rule designer's payoff is strictly higher and therefore it cannot be that $\sigma_\varepsilon = 0$. In principle, it may be optimal to choose an intermediate value for the noise σ_ε to induce the optimizing type to do something better than the static best response.

The results in Proposition 4 are also informative about the optimal tenure of the policy maker. In fact, an alternative instrument for the rule designer to separate the static policy choice from the evolution of the reputation of the policy maker in subsequent periods is to terminate the current policy maker's tenure after one period. This is equivalent to choosing a perfectly opaque rule with $\sigma_\varepsilon = \infty$. Thus early termination (one-period tenure) is optimal when the reputation of a new policy maker is sufficiently high.

5.2 Stochastic rules

An alternative way of introducing opacity in rules is to allow the rule designer to choose stochastic rules even though fundamentals are deterministic. The rule designer can now choose a rule that consists of a set of policies, Σ_c, and a probability distribution over these policies, σ_c. We can interpret this as introducing clauses that allow policies to be conditioned on irrelevant details. The commitment type will then draw a policy from this distribution. The optimizing type can also randomize across policies. We will denote its strategy as σ_o.

The rule designer's problem is then

$$\max_{x, \sigma_o, \sigma_c} \int \left[w(x, \pi) + \beta W_0 \left(\rho'(\pi | \sigma_o, \sigma_c, \rho) \right) \right] \left[\rho \sigma_c(\pi) + (1-\rho) \sigma_o(\pi) \right] d\pi \tag{15}$$

subject to $\sigma_c, \sigma_o \in \Delta([\underline{\pi}, \bar{\pi}])$, the implementability condition,

$$x = \phi\left(\int \pi\left[\rho\sigma_c(\pi) + (1-\rho)\sigma_o(\pi)\right] d\pi\right), \tag{16}$$

and the incentive compatibility constraint for the optimizing type, $\forall \pi \in \mathrm{Supp}\sigma_o, \forall \tilde{\pi} \in \mathrm{Supp}\sigma_o \cup \mathrm{Supp}\sigma_c \cup \{\pi^*(x)\}$

$$w(x, \pi) + \beta V_0\left(\rho'(\pi|\sigma_o, \sigma_c, \rho)\right) \geqslant w(x, \tilde{\pi}) + \beta V_0\left(\rho'(\tilde{\pi}|\sigma_o, \sigma_c, \rho)\right), \tag{17}$$

given the evolution of beliefs,

$$\rho'(\pi|\sigma_o, \sigma_c, \rho) = \frac{\rho\sigma_c(\pi)}{\rho\sigma_c(\pi) + (1-\rho)\sigma_o(\pi)}. \tag{18}$$

We say that a rule is stochastic if the support of σ_c, Σ_c, contains more than one element, while a rule is deterministic if the support of σ_c is a singleton. Similar to Proposition 4, we show that if the policy maker's reputation is high enough, then recommending stochastic rules is optimal, while if its reputation is sufficiently close to zero, then it is optimal to have deterministic rules that provide strong incentives for the optimizing type.

Proposition 5. *Suppose Assumptions 1–4 hold.* Then, there exist ρ_1 and ρ_2 with $0 < \rho_1 \leqslant \rho_2 < 1$ such that:

1. *For* $\rho \in [\rho_2, 1]$ *it is optimal to have stochastic rules.*

2. *For* $\rho \in [0, \rho_1]$ *a deterministic rule is optimal and, in particular,* $\pi_c = \pi_{ico}(\rho)$ *with probability one.*

Consider first the case in which the policy maker's reputation is close to one. The optimality of stochastic rules follows from the properties of Bayes' rule and continuation values being increasing in the prior and does not rely on uncertainty being beneficial. To establish the result, suppose by way of contradiction that it is optimal to choose a rule that recommends policy $\underline{\pi}$ with probability one. This is the best deterministic rule, as shown in Proposition 2. Consider a perturbation in which the rule puts a small but positive probability, ε, on the static best response. When ρ is close to one, on observing the static best response, the agents attribute it to the perturbation of the commitment type rather than to the optimizing type. Consequently, the posterior that the policy maker is the commitment type rises sharply, which increases the expected continuation value of the perturbation and more than compensates the static losses.[14]

[14]Note that forcing the commitment type to randomize reduces the variance of the posterior. In fact, under the deterministic rule with separation, the posterior is one with probability ρ and zero with probability

The case with reputation close to zero instead relies on uncertainty being beneficial. The argument here mirrors the one provided to show that randomization by the optimizing type is not optimal. Randomization tightens the optimizing type's incentive constraint, resulting in a more lenient expected policy, which in turn lowers the static payoff in addition to the dynamic losses that arise because uncertainty is beneficial.

The message of this section is that when the policy maker's reputation is low, rules should be transparent and easily interpretable so that deviations are easily detectable. This is because providing incentives to the optimizing type is critical, as in Atkeson et al. (2007). In contrast, when the policy maker's reputation is high, rules should be opaque and hard to interpret. This is because the benefits of maintaining uncertainty about the policy maker's type outweigh the costs associated with looser incentives to the optimizing type. This interpretation can account for why policy institutions with arguably high credibility (such as the U.S. Federal Reserve) do not rely on strict numerical rules. For instance, in an influential policy speech about the conduct of monetary policy, the Fed chairman Jerome Powell said:

> In seeking to achieve inflation that averages 2 percent over time, we are not tying ourselves to a particular mathematical formula that defines the average. Thus, our approach could be viewed as a flexible form of average inflation targeting. Our decisions about appropriate monetary policy will continue to reflect a broad array of considerations and will not be dictated by any formula. (Powell, 2020)

Moreover, policy rules are often based on soft numerical targets that leave room for multiple interpretations and are hard to monitor. For example, chairman Powell said "specifying a numerical goal for employment is unwise, because the maximum level of employment is not directly measurable and changes over time for reasons unrelated to monetary policy (Powell, 2020)." Similarly, the fiscal rules in the European monetary union rely on arguably opaque measures of output gaps provided by the European Commission itself.[15]

6 Role of private agents

In this section, we show that the rule designer's desire to influence the actions of the private agents is critical for our results. The macroeconomics literature on rules and time inconsistency typically considers two types of models. First, in the tradition of Kydland

$1-\rho$. Under our perturbation, the posterior is one with probability $\rho(1-\varepsilon)$ and $\rho\varepsilon/[\rho\varepsilon+(1-\rho)]$ with probability $\rho\varepsilon+(1-\rho)$.

[15] See https://www.europarl.europa.eu/RegData/etudes/BRIE/2016/574407/IPOL_BRI(2016)574407_EN.pdf.

and Prescott (1977), there are models in which the policy maker is benevolent (i.e., has the same preferences as society) and time-inconsistency problems arise because of externalities associated with the presence of private agents; see Chari et al. (1988). Examples include Barro and Gordon (1983b), Chari and Kehoe (1990), Athey et al. (2005), and Kareken and Wallace (1978). Our setup follows in this tradition. Second, the literature also considers principal–agent or delegation economies in which the policy maker (the agent) has different preferences than society (the principal). This difference results in a conflict of interest between the principal and the agent which generates a time inconsistency problem. Examples include Amador et al. (2006), Amador and Bagwell (2013), and Halac and Yared (2014).

One implication of the principal–agent framework is that under commitment, if the agent has any private information, the principal always prefers to learn the agent's private information. In other words, uncertainty is *never* beneficial. The critical difference between our model and the principal–agent framework is the presence of a third agent, whom we call private agents, that take an action otherwise taken by the principal.[16] It is because of the rule designer's desire to influence the actions of these private agents that uncertainty may be beneficial. Next, we illustrate this claim by means of a simple example.

Consider a principal–agent or delegation framework. There is a rule designer (principal) and a policy maker (agent). To better relate to this literature, we consider the preference type interpretation of the policy maker discussed in Section 2. The policy maker can be one of two types, $\theta \in \{\theta_o, \theta_c\}$, where θ_c corresponds to the commitment type and θ_o to the optimizing type. The type θ determines the cost of deviating from the policy recommended by the rule designer.[17] Define ρ to be the rule designer's prior of facing the θ_o type. At the beginning of the period, the rule designer takes an action x and proposes a menu of actions $\pi_r(\theta)$ to the policy maker. The policy maker (agent) then takes an action π. The rule designer's preferences are

$$w(x, \pi),$$

and the policy maker's preferences are

$$v(x, \pi) - \theta \left|\pi - \pi_r(\theta)\right|.$$

[16] This difference in preferences between the policy maker and the rule designer is not critical. For example, it is easy to show that our conclusions hold in a Barro–Gordon model where the policy maker has a different inflation bias than society.

[17] Alternatively, we can assume that the rule designer can impose a punishment or reward $\zeta \in [0, \bar{\zeta}]$ on the agent after it observes the agent's action and θ measures the value of this punishment or reward to the agent.

We assume that θ_c is sufficiently large so that the rule designer can effectively control the action of the θ_c type agent and that $\theta_o = 0$. Thus, as in our basic framework, the rule designer directly chooses the action of the θ_c (commitment) type but not of the θ_o (optimizing) type. The principal here chooses x directly without having to influence the private agents to choose a desirable level of x. This is the critical difference from our environment.

We next show that in this principal–agent setting, uncertainty is never beneficial and pooling is never optimal.

Define $\pi^*_v(x) \equiv \arg\max_\pi v(x, \pi)$. Consider first the static principal–agent problem:

$$W_0(\rho) = \max_{x, \pi_c, \pi_o} \rho w(x, \pi_c) + (1 - \rho) w(x, \pi_o) \tag{19}$$

subject to the incentive constraint for the optimizing type

$$\pi_o = \pi^*_v(x).$$

Note that the problem (19) differs from the static problem in our economy, (3), because there is no implementability constraint (4).

Proposition 6. *In the principal–agent economy, uncertainty is never beneficial in that $W_0(\rho) \leqslant \rho W(1) + (1 - \rho) W_0(0)$.*

The proof for this proposition is straightforward. For any $\rho \in (0, 1)$ we have

$$\begin{aligned} W_0(\rho) &= \max_{x, \pi_c} \{\rho w(x, \pi_c) + (1 - \rho) w(x, \pi^*(x))\} \\ &\leqslant \rho \max_{x, \pi_c} w(x, \pi_c) + (1 - \rho) \max_x w(x, \pi^*(x)) \\ &= \rho W_0(1) + (1 - \rho) W_0(0). \end{aligned}$$

The idea is simply that with more information the rule designer can better tailor x to the policy π chosen by the policy maker. When the rule designer knows the policy maker's type, it can condition x to the type, while when it is uncertain it must choose a single value of x. In our environment instead, the rule designer does not control x directly and must influence the private agents to take desirable actions. Under the sufficient conditions described in Section 3, learning the policy maker's type can incentivize the private agents to choose a very unfavorable action from the principal's perspective. Therefore, the rule designer prefers if the private agents do not learn the policy maker's type. This leads the rule designer to choose a policy recommendation which does not lead to the revelation of the policy maker's type.

Consider now the twice repeated problem for the principal–agent economy. When the rule designer has commitment, we can think of its problem as choosing an action x, proposing policies π_c and π_o, and promised values for the optimizing type conditional on the chosen policy, $V_o(\pi)$. Without loss of generality, we can set $V_o(\pi) = \underline{V}$ for all $\pi \notin \{\pi_o, \pi_c\}$, where the value $\underline{V} \equiv \min_x \max_\pi v(x, \pi)$ is the worst continuation value that can be promised to the optimizing type. The rule designer's problem is

$$\bar{W}(\rho) = \max_{x, \pi_c, \pi_o, V_o(\pi)} \rho \left[w(x, \pi_c) + \beta \bar{W}_0 \left(V_o(\pi_c), \rho_c' \right) \right] + (1 - \rho) \left[w(x, \pi_o) + \beta \bar{W}_0 \left(V_o(\pi_o), \rho_o' \right) \right]$$

subject to the incentive constraints for the θ_o type,

$$v(x, \pi_o) + \beta V_o(\pi_o) \geqslant v(x, \pi) + \beta \underline{V} \quad \forall \pi \notin \{\pi_o, \pi_c\},$$

$$v(x, \pi_o) + \beta V_o(\pi_o) \geqslant v(x, \pi_c) + \beta V_o(\pi_c).$$

The two incentive constraints require that the θ_o type prefers to follow its recommended action rather than engage in either a detectable deviation or follow the recommended action for the θ_c type respectively. The rule designer's continuation value $\bar{W}_0$ given promised value V and prior ρ is

$$\bar{W}_0(V, \rho) = \max_{x, \pi_c, \pi_o} \rho w(x, \pi_c) + (1 - \rho) w(x, \pi_o)$$

subject to the incentive constraint for the optimizing type,

$$v(x, \pi_o) \geqslant v(x, \pi) \quad \forall \pi,$$

and the optimizing type's promise-keeping constraint,

$$v(x, \pi_o) = V.$$

Proposition 7. *In the twice repeated principal–agent economy, pooling is never optimal.*

The proof of this proposition follows from two observations. First, similar to Proposition 6, uncertainty is not beneficial in the last period because by knowing the policy maker's type, the principal can better tailor its action to the policy. This implies that pooling does not have dynamic benefits, as defined in Section 4.1. The second observation is that providing incentives to the optimizing type in the first period does not require that there be no revelation of information, and thus the preservation of the optimizing type's reputation. This is because the rule designer can promise continuation value $V_o > \underline{V}$ by committing to an action x in the second period that is more favorable to the optimizing type conditional on the optimizing type following the recommendation. This implies that pooling also does not have static benefits, as defined in Section 4.1.[18] In our environment instead, the rule designer cannot force private agents to reward the optimizing type in the second period if the rule designer loses its reputation. Therefore, the loss in reputation is the only way the rule designer can provide incentives to the optimizing type even if the rule designer has commitment (See Appendix E.1).

7 Conclusion

In this paper, we study the optimal design of rules in a dynamic model when there is a time inconsistency problem and uncertainty about whether the policy maker can commit to follow the rule ex post. We show that in a large class of economies preserving uncertainty about the policy maker's type is preferable from an ex-ante perspective. Therefore, learning the type of the policy maker can be costly. When the policy maker's reputation is low, we show that reputational considerations imply that the optimal rule is more lenient than the one that would arise in a static environment. For example, in the context of financial regulation, on-path bailouts are necessary to discipline future risk taking by financial institutions. Moreover, opaque rules are preferable to transparent ones when reputation is high.

In our analysis we abstract from the question on the optimal-degree flexibility when the policy maker has private information about the state of the economy considered by the delegation literature; see, for instance Athey et al. (2005), Amador et al. (2006), and Halac and Yared (2014). In our economy with no fundamental uncertainty, if the policy maker follows the rule for sure, it is trivially optimal to leave no flexibility, and by doing so, implement the Ramsey outcome. An interesting avenue for future work is to study how the incentives to build reputation considered in this paper interact with the choice of how much flexibility to leave to the policy maker.

References

AMADOR, M. AND K. BAGWELL (2013): "The theory of optimal delegation with an application to tariff caps," *Econometrica*, 81, 1541–1599. 6, 28

AMADOR, M. AND C. PHELAN (2018): "Reputation and Sovereign Default," Tech. rep., National Bureau of Economic Research. 6

[18]If the principal lacks commitment, then pooling may be a costly way to commit to reward the optimizing type ex post.

Amador, M., I. Werning, and G.-M. Angeletos (2006): "Commitment vs. flexibility," *Econometrica*, 74, 365–396. 6, 28, 31

Angeletos, G.-M., C. Hellwig, and A. Pavan (2006): "Signaling in a global game: Coordination and policy traps," *Journal of Political economy*, 114, 452–484. 7

Asriyan, V., D. Foarta, and V. Vanasco (2019): "The good, the bad and the complex: Product design with imperfect information," Tech. rep. 7

Athey, S., A. Atkeson, and P. J. Kehoe (2005): "The optimal degree of discretion in monetary policy," *Econometrica*, 73, 1431–1475. 6, 28, 31

Atkeson, A., V. V. Chari, and P. J. Kehoe (2007): "On the optimal choice of a monetary policy instrument," . 4, 23, 27

Atkeson, A. and P. J. Kehoe (2001): "The advantage of transparent instruments of monetary policy," Tech. rep., National Bureau of Economic Research. 4, 23

Azzimonti, M., M. Battaglini, and S. Coate (2016): "The costs and benefits of balanced budget rules: Lessons from a political economy model of fiscal policy," *Journal of Public Economics*, 136, 45–61. 6

Backus, D. and J. Driffill (1985): "Inflation and reputation," *American economic review*, 75, 530–538. 6

Barro, R. J. (1986): "Reputation in a model of monetary policy with incomplete information," *Journal of Monetary Economics*, 17, 3–20. 3, 6

Barro, R. J. and D. B. Gordon (1983a): "A positive theory of monetary policy in a natural rate model," *Journal of political economy*, 91, 589–610. 7, 9

——— (1983b): "Rules, discretion and reputation in a model of monetary policy," *Journal of monetary economics*, 12, 101–121. 3, 6, 28

Bond, P. and Y. Zeng (2018): "Silence is Safest: Non-Disclosure When the Audience's Preferences are Uncertain," *Available at SSRN 3124106*. 7

Chari, V. V. and P. J. Kehoe (1990): "Sustainable plans," *Journal of political economy*, 98, 783–802. 6, 7, 28

Chari, V. V., P. J. Kehoe, and E. C. Prescott (1988): "Time consistency and policy," Tech. rep., Federal Reserve Bank of Minneapolis. 6, 28

COLE, H. L., J. DOW, AND W. B. ENGLISH (1995): "Default, settlement, and signalling: Lending resumption in a reputational model of sovereign debt," *International Economic Review*, 365–385. 7

CRÉMER, J. (1995): "Arm's length relationships," *The Quarterly Journal of Economics*, 110, 275–295. 5

DEBORTOLI, D. AND R. NUNES (2010): "Fiscal policy under loose commitment," *Journal of Economic Theory*, 145, 1005–1032. 7

DOVIS, A. AND R. KIRPALANI (2020a): "Fiscal rules, bailouts, and reputation in federal governments," *American Economic Review*, 110, 860–88. 6, 7, 77

——— (2020b): "Reputation, Bailouts, and Interest Rate Spread Dynamics," . 6

HALAC, M. AND P. YARED (2014): "Fiscal rules and discretion under persistent shocks," *Econometrica*, 82, 1557–1614. 6, 28, 31

——— (2018a): "Fiscal rules and discretion in a world economy," *American Economic Review*, 108, 2305–34. 6

——— (2018b): "Fiscal Rules and Discretion under Self-Enforcement," . 6

KAREKEN, J. H. AND N. WALLACE (1978): "Deposit insurance and bank regulation: A partial-equilibrium exposition," *Journal of Business*, 413–438. 3, 6, 7, 9, 28

KING, R. G., Y. K. LU, AND E. S. PASTEN (2008): "Managing expectations," *Journal of Money, Credit and banking*, 40, 1625–1666. 7

KREPS, D. M. AND R. WILSON (1982): "Reputation and imperfect information," *Journal of economic theory*, 27, 253–279. 6

KYDLAND, F. E. AND E. C. PRESCOTT (1977): "Rules rather than discretion: The inconsistency of optimal plans," *Journal of political economy*, 85, 473–491. 2, 6, 27

LU, Y. K. (2013): "Optimal policy with credibility concerns," *Journal of Economic Theory*, 148, 2007–2032. 7

LU, Y. K., R. G. KING, AND E. PASTEN (2016): "Optimal reputation building in the New Keynesian model," *Journal of Monetary Economics*, 84, 233–249. 7

MARINOVIC, I. AND M. SZYDLOWSKI (2019): "Monitor Reputation and Transparency," . 7

MILGROM, P. AND J. ROBERTS (1982): "Predation, reputation, and entry deterrence," *Journal of economic theory*, 27, 280–312. 6

NOSAL, J. B. AND G. ORDOÑEZ (2016): "Uncertainty as commitment," *Journal of Monetary Economics*, 80, 124–140. 3

PHELAN, C. (2006): "Public trust and government betrayal," *Journal of Economic Theory*, 130, 27–43. 6

PHELAN, C. AND E. STACCHETTI (2001): "Sequential equilibria in a Ramsey tax model," *Econometrica*, 69, 1491–1518. 6

PIGUILLEM, F. AND A. RIBONI (2018): "Fiscal Rules as Bargaining Chips," Tech. rep. 5

PIGUILLEM, F. AND A. SCHNEIDER (2013): "Coordination, Efficiency and Policy Discretion," Tech. rep., Einaudi Institute for Economics and Finance (EIEF). 4, 23

POWELL, J. H. (2020): "New Economic Challenges and the Fed's Monetary Policy Review: a speech at" Navigating the Decade Ahead: Implications for Monetary Policy," an economic policy symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 27, 2020," Tech. rep. 27

PRAT, A. (2005): "The Wrong Kind of Transparency," *The American Economic Review*, 95, 862–877. 5

ROGOFF, K. (1985): "The optimal degree of commitment to an intermediate monetary target," *The quarterly journal of economics*, 100, 1169–1189. 76, 80

SANKTJOHANSER, A. (2018): "Optimally Stubborn," Tech. rep., working paper. 7

VICKERS, J. (1986): "Signalling in a model of monetary policy with incomplete information," *Oxford Economic Papers*, 38, 443–455. 7

Appendix (Not for publication)

A Omitted proofs

A.1 Proof of Proposition 1

Here we provide the proof for the case in which $w_x < 0$. The case with $w_x > 0$ follows from a symmetric argument.

Suppose first that $\pi_c(\rho) = \underline{\pi}$ for all ρ. We first prove a preliminary result.

Claim. Under Assumptions 1 and 2, if $w_x < 0$ then $\pi^*(x)$ is increasing and convex and $x(\rho)$ is decreasing and convex.

Proof. The optimizing type's static best response, $\pi^*(x)$, is implicitly defined by the first order condition:

$$w_\pi(x, \pi^*(x)) = 0 \tag{20}$$

so by the implicit function theorem we have

$$\pi^*_x(x) = -\frac{w_{x\pi}(x, \pi^*(x))}{w_{\pi\pi}(x, \pi^*(x))} > 0 \tag{21}$$

where the inequality follows from $w_{x\pi} > 0$ and $w_{\pi\pi} < 0$.

Implicitly differentiating (20) twice we have

$$\pi^*_{xx}(x) = \frac{w_{\pi\pi\pi}(x, \pi^*(x))\,\pi^*_x(x)^2 + 2w_{\pi\pi x}\pi^*_x(x) + w_{x\pi x}(x, \pi^*(x))}{(-w_{\pi\pi}(x, \pi^*(x)))}$$

Thus, in order for $\pi^*(x)$ to be convex we need

$$w_{\pi\pi\pi}(x, \pi^*(x))\,\pi^*_x(x)^2 + 2w_{\pi\pi x}\pi^*_x(x) + w_{x\pi x}(x, \pi^*(x)) \geqslant 0.$$

Notice that the expression above can be equivalently written as

$$[\pi^*_x(x), 1]\, \nabla^2 w_\pi(x, \pi) \begin{bmatrix} \pi^*_x(x) \\ 1 \end{bmatrix}$$

which is positive since $w_\pi(x, \pi)$ is convex by Assumption 2. Hence $\pi^*(x)$ is increasing and convex.

Consider now

$$x(\rho) = \phi(\rho\pi_c + (1-\rho)\,\pi_o(\rho)) \tag{22}$$

where $\pi_o(\rho) = \pi^*(x(\rho))$. So

$$x'(\rho) = \frac{\phi'[\pi_c - \pi_o]}{[1 - \phi'(1-\rho)\pi_x^*(x(\rho))]} < 0 \tag{23}$$

which is negative since $\pi_o \geqslant \underline{\pi}$ and $1 - \phi'(1-\rho)\pi_x^*(x) \geqslant 0$, where the latter follows from condition 3 of Assumption 2. Twice differentiating the implementability condition (22) we obtain

$$\begin{aligned} x''(\rho) &= \phi''\left[\pi_c - \pi_o + (1-\rho)\pi_o'(\rho)\right]^2 + \phi'\left[-\pi_o'(\rho) + (1-\rho)\pi_o''(\rho) - \pi_o'(\rho)\right] \\ &= \phi''\left[\pi_c - \pi_o + (1-\rho)\pi_o'(\rho)\right]^2 + \phi'\left[(1-\rho)\pi_o''(\rho) - 2\pi_x^*(x)\,x'(\rho)\right] \end{aligned} \tag{24}$$

and since $\pi_o(\rho) = \pi^*(x(\rho))$ then

$$\pi_o''(\rho) = \pi_{xx}^*(x)\,x'(\rho)^2 + \pi_x^*(x)\,x''(\rho). \tag{25}$$

Therefore, using (25) to substitute for $\pi_o''(\rho)$ in (24), we obtain

$$x''(\rho) = \frac{\phi''[\pi_c - \pi_o + (1-\rho)\pi_o'(\rho)]^2 + \left[\phi'(1-\rho)\pi_{xx}^*(x)\,x'(\rho)^2 - 2\phi'\pi_x^*(x)\,x'(\rho)\right]}{[1 - \phi'(1-\rho)\pi_x^*(x)]} \tag{26}$$

Thus $x'' \geqslant 0$ follows from Assumption 1, $\phi' \geqslant 0$, $\phi'' \geqslant 0$, $\pi_x^*(x) > 0$, and $\pi_{xx}^*(x) \geqslant 0$, where the last two inequalities were proved earlier. □

We now turn the proof of the Proposition. Define

$$\bar{w}(x) = w(x, \underline{\pi})$$

$$w^*(x) = w(x, \pi^*(x))$$

and

$$F(\rho, x(\rho)) = \rho\bar{w}(x(\rho)) + (1-\rho)\,w^*(x(\rho))$$

and so $W_0(\rho) = F(\rho, x(\rho))$. We want to show that $W_0(\rho)$ is concave. We have,

$$\begin{aligned} W_0''(\rho) &= \left[1, x'(\rho)\right]\nabla^2 F(\rho, x(\rho))\begin{bmatrix} 1 \\ x'(\rho)\end{bmatrix} + \nabla F(\rho, x(\rho))\begin{bmatrix} 0 \\ x''(\rho)\end{bmatrix} \\ &= 2F_{\rho x}x'(\rho) + F_{xx}x'(\rho)^2 + F_x(\rho, x(\rho))\,x''(\rho). \end{aligned} \tag{27}$$

Thus to prove the result it is sufficient to show that the above expression is negative. We proceed in several steps. First we show that $F_{xx}x'(\rho)^2 < 0$ and then we show that

$2F_{\rho x}x'(\rho) + F_x(\rho, x(\rho))\, x''(\rho) \leqslant 0.$

To see that $F_{xx}x'(\rho)^2 < 0$, note that

$$F_x = \rho w_x(x, \pi_c) + (1-\rho)\, w_x(x, \pi_o) < 0 \tag{28}$$

since $w_{x\pi} > 0$, and so

$$\begin{aligned} F_{xx} &= \rho \bar{w}_{xx}(x(\rho)) + (1-\rho)\, w^*_{xx}(x(\rho)) \\ &= \rho w_{xx}(x, \pi_c) + (1-\rho)\left[w_{xx}(x, \pi_o(x)) + w_{x\pi}(x, \pi)\, \pi^*_x(x)\right]. \end{aligned} \tag{29}$$

The first term in (29) is negative since $w_{xx} < 0$. For the second, note that from (21), we have that

$$\pi^*_x(x) = -\frac{w_{x\pi}(x, \pi_o)}{w_{\pi\pi}(x, \pi_o)} \tag{30}$$

Therefore, using (30) we can write the second term in (29) as

$$w_{xx}(x, \pi_o) + w_{x\pi}(x, \pi_o)\, \pi^*_x(x) = \frac{w_{xx}(x, \pi_o)\, w_{\pi\pi}(x, \pi_o) - w_{x\pi}(x, \pi_o)^2}{w_{\pi\pi}(x, \pi_o)} < 0$$

where the inequality follows from w being concave in (x, π). Therefore, both terms in (29) are negative and so $F_{xx}x'(\rho)^2 < 0$.

Now note that we can write $F_{\rho x}$ as

$$\begin{aligned} F_{\rho x} &= \bar{w}_x(x(\rho)) - w^*_x(x(\rho)) = w_x(x, \pi_c) - \left[w_x(x, \pi^*(x)) + w_\pi(x, \pi^*(x))\, \pi^*_x(x)\right] \\ &= w_x(x, \pi_c) - w_x(x, \pi^*(x)) < 0. \end{aligned} \tag{31}$$

Then from (27) it follows that

$$\begin{aligned} W_0''(\rho) =& 2F_{\rho x}x'(\rho) + F_{xx}x'(\rho)^2 + F_x(\rho, x(\rho))\, x''(\rho) \\ <& 2F_{\rho x}x'(\rho) + F_x(\rho, x(\rho))\, x''(\rho) \\ =& 2\left[w_x(x, \pi_c) - w_x(x, \pi^*(x))\right] x'(\rho) + \left[\rho w_x(x, \pi_c) + (1-\rho)\, w_x(x, \pi^*(x))\right] x''(\rho) \end{aligned}$$

where the first inequality follows from the fact that $F_{xx}x'(\rho)^2 < 0$ and the second equality follows from (31) and (28). A sufficient condition for the concavity of W_0 is then

$$2\left[w_x(x, \pi_c) - w_x(x, \pi^*(x))\right] x'(\rho) + \left[\rho w_x(x, \pi_c) + (1-\rho)\, w_x(x, \pi^*(x))\right] x''(\rho) \leqslant 0$$

or, rearranging terms,

$$\frac{2\left[w_x(x, \pi^*(x)) - w_x(x, \pi_c)\right]}{-\left[\rho w_x(x, \pi_c) + (1-\rho)\, w_x(x, \pi^*(x))\right]} \leqslant -\frac{x''(\rho)}{x'(\rho)}. \tag{32}$$

Using the expressions for $x'(\rho)$ and $x''(\rho)$ in (23) and (26) we have

$$
\begin{aligned}
-\frac{x''(\rho)}{x'(\rho)} &= \frac{\phi''\left[\pi_c - \pi_o + (1-\rho)\,\pi_o'(\rho)\right]^2 + \left[\phi'(1-\rho)\,\pi_{xx}^*(x)\,x'(\rho)^2 - 2\phi'\pi_x^*(x)\,x'(\rho)\right]}{\phi'\left[\pi_o - \pi_c\right]} \\
&\geqslant \frac{-2\phi'\pi_x^*(x)}{\phi'\left[\pi_o - \pi_c\right]}x'(\rho) \\
&= \frac{2\phi'\pi_x^*(x)}{\left[1 - \phi'(1-\rho)\,\pi_x^*(x(\rho))\right]}
\end{aligned}
$$

where the first inequality follows from the fact that $\phi'' > 0$, $\phi' > 0$, $\pi_{xx}^* > 0$ and the denominator is positive, and the last equality follows from using (23) to substitute for $x'(\rho)$. Hence a sufficient condition for (32) is

$$
\frac{2\left[w_x(x,\pi^*(x)) - w_x(x,\pi_c)\right]}{-\left[\rho w_x(x,\pi_c) + (1-\rho)\,w_x(x,\pi^*(x))\right]} \leqslant \frac{2\pi_x^*(x)\,\phi'}{\left[1 - \phi'(1-\rho)\,\pi_x^*(x)\right]}
$$

or, rearranging terms,

$$
\pi_x^*(x)\,\phi' + \frac{w_x(x,\pi^*(x))}{w_x(x,\pi_c)} \geqslant 1.
$$

Notice that

$$
\pi_x^*(x)\,\phi' + \frac{w_x(x,\pi^*(x))}{w_x(x,\pi_c)} \geqslant \pi_x^*(\underline{x})\,\phi'(\underline{\pi}) + \frac{w_x(\underline{x},\pi^*(\underline{x}))}{w_x(\underline{x},\underline{\pi})} \geqslant 1
$$

where we have used the fact that $\pi^*(x)\,,\phi(\pi)$ are convex, $w(x,\pi)$ is concave, and the last inequality follows from condition 3 of Assumption 2. Thus $W_0''(\rho) < 0$.

As a final step we will show that our assumptions imply that π_c is independent of ρ and in particular equals $\underline{\pi}$. The first order condition of the static government's problem with respect to π_c is

$$
\rho w_\pi(x,\pi) + \left[\rho w_x(x,\pi) + (1-\rho)\,w_x(x,\pi^*(x))\right]\frac{\rho\phi'(\cdot)}{\left[1 - \phi'(\cdot)(1-\rho)\,\pi_x^*(x)\right]}
$$

Since by assumption

$$
w_\pi(x,\pi) + \left[\rho w_x(x,\pi) + (1-\rho)\,w_x(x,\pi^*(x))\right]\frac{\phi'(\cdot)}{\left[1 - \phi'(\cdot)(1-\rho)\,\pi_x^*(x)\right]} \leqslant 0
$$

it must be that $\pi_c = \underline{\pi}$. Q.E.D.

We next show that the two examples satisfy Assumption 2.

Lemma 1. *The Barro-Gordon economy satisfies Assumption 2.*

Proof. Recall that

$$w(x,\pi) - \frac{1}{2}\left[(\psi + x - \pi)^2 + \pi^2\right]$$

Thus,

$$w_{xx} = -1 < 0$$

$$w_{\pi\pi} = -2 < 0$$

and

$$w_{xx}w_{\pi\pi} - w_{x\pi} = 1 > 0.$$

Therefore the Hessian of $w(x,\pi)$ is negative semi-definite and thus $w(x,\pi)$ is concave.

Next, note that $w_{\pi\pi a} = 0, w_{\pi\pi\pi} = 0$ and $w_{\pi aa} = 0$ and so w_π is convex. Clearly, since $\phi(\pi) = \pi$ then $\phi'' = 0$. Thus, strategic complementarities are decreasing in reputation (Condition 2).

Finally, using a little algebra Condition 3 is

$$\left[-2 + \frac{2\rho}{1+\rho}\right]\pi_c$$

which is strictly less than zero for any π_c positive. Thus, it must be $\pi_c = 0$ and Condition 3 is satisfied. Next, since

$$\pi^*(x) = \frac{x+\psi}{2}$$

we have

$$\pi^*_x(x)\,\phi'(\pi) = \frac{1}{2}.$$

Therefore

$$1 > \pi^*_x(x)\,\phi'(\pi) = \frac{1}{2} \geqslant 1 - \frac{w_x(\underline{x},\pi^*(\underline{x}))}{w_x(\underline{x},\underline{\pi})} = \frac{1}{2}$$

and so Condition 4 is satisfied. Q.E.D.

Lemma 2. *If ψ is sufficiently small, then bailout economy satisfies Assumption 2.*

Proof. Recall that

$$w(e,\pi) = -\nu(e) + p(e)R_H - (1-p(e))(1-\pi)\psi - c(\pi)$$

and thus $\pi^*(e)$ is the solution to

$$(1-p(e))\psi - c'(\pi) = 0.$$

Let's first show that $w(e,\pi)$ is concave. We have

$$w_e = -v'(e) + p'(e)(R_H + (1-\pi)\psi),$$

$$w_{ee} = -v''(e) + p''(e)(R_H + (1-\pi)\psi) < 0,$$

$$w_{e\pi} = -p''(e)\psi > 0,$$

$$w_\pi = (1-p(e))\psi - c'(\pi),$$

$$w_{\pi\pi} = -c''(\pi) \leqslant 0.$$

So

$$\begin{aligned} & w_{ee}w_{\pi\pi} - w_{e\pi}^2 \\ = & \left[-v''(e) + p''(e)(R_H + (1-\pi)\psi)\right]\left(-c''(\pi)\right) - p''(e)^2\psi^2. \end{aligned} \tag{33}$$

The first term is positive since $v'' > 0$, $p'' < 0$, and $c'' > 0$ but $-p''(e)^2\psi^2$ is negative. Clearly, the whole expression is positive if ψ is small enough. Thus the Hessian of w is negative semi-definite which implies that w is concave.

Next, we show that $w_\pi(e,\pi)$ is convex. We have

$$w_\pi(e,\pi) = (1-p(e))\psi - c'(\pi).$$

Therefore

$$w_{\pi e} = -p'(e)\psi,$$

$$w_{\pi ee} = -p''(e)\psi > 0,$$

$$w_{\pi e\pi} = 0,$$

$$w_{\pi\pi\pi} = -c'''(\pi) = 0.$$

since $c(e)$ is quadratic and so $c''' = 0$. Therefore

$$w_{\pi ee}w_{\pi\pi\pi} - w_{\pi e\pi}^2 = \left[-p'(e)\psi\right]\left[-c'''(e)\right] = 0.$$

Thus the Hessian of w_π is positive semi-definite and so w_π is convex. Under our functional form assumptions,

$$p(e) = e^\alpha$$

$$c(\pi) = \lambda\pi^2/2$$

$$v(e) = e^2/2,$$

we have that

$$e = \phi(\pi) = \alpha^{\eta}(R_H - \pi)^{\eta}, \quad \eta \equiv 1/(2-\alpha) \in (0,1)$$

so ϕ is decreasing and concave. Thus, strategic complementarities are decreasing in reputation (Condition 2).

We now check that Condition 4 is satisfied. We have to show that the following two conditions hold:

$$1 > \pi_x^*(x)\,\phi'(\pi),$$
$$\pi_x^*(x)\,\phi'(\pi) \geqslant 1 - \frac{w_x(\underline{x}, \pi^*(\underline{x}))}{w_x(\underline{x}, \underline{\pi})}.$$

Notice that

$$\pi_x^*(x)\,\phi'(\pi) = \left(-\frac{p'(e)\,\psi}{c''(\pi)}\right)\phi'(\pi)$$

and

$$1 - \frac{w_x(\underline{x}, \pi^*(\underline{x}))}{w_x(\underline{x}, \underline{\pi})} = 1 - \frac{-\nu'(\underline{e}) + p'(\underline{e})\,R_H + p'(\underline{e})\,(1-\pi^*(\underline{e}))\,\psi}{-\nu'(\underline{e}) + p'(\underline{e})\,R_H + p'(\underline{e})\,\psi} = \pi^*(\underline{e}).$$

Thus, the two conditions can be written as

$$1 > \left(-\frac{p'(e)\,\psi}{c''(\pi)}\right)\phi'(\pi) \tag{34}$$

and

$$\left(-\frac{p'(e)\,\psi}{c''(\pi)}\right)\phi'(\pi) \geqslant \pi^*(\underline{e}). \tag{35}$$

Under our functional form assumptions,

$$\begin{aligned}\left(-\frac{p'(e)\,\psi}{c''(\pi)}\right)\phi'(\pi) &= \frac{-\alpha e^{\alpha-1}\psi}{\lambda}\eta\alpha^{\eta}(R_H-\pi)^{\eta-1}\\ &= \frac{-\alpha\left(\alpha^{\eta}(R_H-\pi)^{\eta}\right)^{\alpha-1}\psi}{\lambda}\left(-\eta\alpha^{\eta}(R_H-\pi)^{\eta-1}\right)\\ &= \frac{\alpha^{1+\eta\alpha}}{2-\alpha}\frac{1}{(R_H-\pi)^{1-\eta\alpha}}\frac{\psi}{\lambda}\end{aligned}$$

and since $\underline{e} = \alpha^{\eta}R_H^{\eta}$ we have that

$$\pi^*(\underline{e}) = \frac{1 - p\left(\alpha^{\eta}R_H^{\eta}\right)}{\lambda}\psi = \left[1 - (\alpha R_H)^{\eta\alpha}\right]\frac{\psi}{\lambda}.$$

By inspection, the first inequality, (34), is satisfied if ψ is sufficiently small while the second inequality, (35), is satisfied if R_H is sufficiently large. In fact, as $R_H \to 1/\alpha$, $p(e) \to 1$

so $\pi^*(\underline{e}) \to 0$ while $\left(-\frac{p'(e)\psi}{c''(\pi)}\right)\phi'(\pi) > 0$.

Finally, we check that Condition 3 holds. We have

$$w_\pi(e,\pi) + [\rho w_e(e,\pi) + (1-\rho)w_e(e,\pi^*(e))]\frac{\phi'(\cdot)}{[1-\phi'(\cdot)(1-\rho)\pi_e^*(e)]}$$

$$= (1-p(e))\psi - c'(\pi)$$

$$+\left[\rho p'(e)[\pi + (1-\pi)\psi] + (1-\rho)[\pi_o(e) + (1-\pi^*(e))\psi]\right]\frac{\phi'(\cdot)}{\left[1-\phi'(\cdot)(1-\rho)\left(-\frac{p'(e)\psi}{c''(\pi)}\right)\right]}$$

which is negative if ψ is sufficiently small since $c' > 0, w_e \geqslant 0$, $\phi' \leqslant 0$, and $1-\phi'\pi_e^* \geqslant 0$. Q.E.D.

A.2 Proof of Proposition 2

Consider first ρ close to 1. Since the incentive compatibility is binding, we have that for some $\delta > 0$, for all ρ

$$W_{ramsey} - \delta \geqslant W_{pool}(\rho).$$

Clearly, at $\rho = 1$, $W_{sep}(1)$ attains the Ramsey outcome. By continuity, there exists a ε_δ sufficiently small such that for all $\rho \in (1-\varepsilon_\delta, 1)$,

$$W_{sep}(\rho) \geqslant W_{ramsey} - \delta.$$

Combining the two expressions above we have that for all $\rho \in (1-\varepsilon_\delta, 1)$,

$$W_{pool}(\rho) < W_{sep}(\rho)$$

as wanted.

Consider now ρ close to zero and assume that Assumptions 1 and 2 hold. Thus by Proposition 1 uncertainty is beneficial, $W_0(\rho) > \rho W_0(1) + (1-\rho)W_0(0)$, and the continuation value is higher under pooling than under the separation policy. To show that it is optimal to pool, it is sufficient to show that the static benefits of pooling are positive for priors close to zero, i.e. $\Delta\omega(\rho) \geqslant 0$ where

$$\Delta\omega(\rho) = w(\phi(\pi_{ico}(\rho)), \pi_{ico}(\rho)) - W_0(\rho)$$

and

$$W_0(\rho) = [\rho w(\phi(\rho\pi_c + (1-\rho)\pi_o), \pi_c) + (1-\rho)w(\phi(\rho\pi_c + (1-\rho)\pi_o), \pi_o)].$$

To this end, note that at $\rho = 0$ we have $\Delta\omega(0) = 0$ since $\pi_{ico}(0) = \pi_o(0) = \pi^*(\phi(\pi_o(0)))$. Thus it is sufficient to show that $\Delta\omega'(0) > 0$. Note that

$$\begin{aligned}W_0'(\rho) &= w(\phi(\rho\pi_c + (1-\rho)\pi_o), \pi_c) - w(\phi(\rho\pi_c + (1-\rho)\pi_o), \pi_o) \\ &\quad + [\rho w_x(\phi(\rho\pi_c + (1-\rho)\pi_o), \pi_c) + (1-\rho) w_x(\phi(\rho\pi_c + (1-\rho)\pi_o), \pi_c)]\, \phi'(\rho\pi_c + (1-\rho)\pi_o) \\ &\quad \times \left[\pi_c - \pi_o + (1-\rho)\pi_o'(\rho)\right]\end{aligned}$$

where we used that $w_\pi(\phi(\rho\pi_c + (1-\rho)\pi_o), \pi_o) = 0$. Therefore,

$$\begin{aligned}\Delta\omega'(0) &= \left[w_x(\phi(\pi_{ico}(0)), \pi_{ico}(0))\, \phi'(\pi_{ico}(0)) + w_\pi(\phi(\pi_{ico}(0)), \pi_{ico}(0))\right] \pi_{ico}'(0) \\ &\quad - \{w(\phi(\pi_o(0)), \pi_c) - w(\phi(\pi_o(0)), \pi_o(0)) \\ &\quad + w_x(\phi(\pi_o(0)), \pi_o(0))\, \phi'(\pi_o(0)) \left[\pi_c - \pi_o + \pi_o'(0)\right]\} \\ &\geqslant \left[w_x(\phi(\pi_{ico}(0)), \pi_{ico}(0))\, \phi'(\pi_{ico}(0)) + w_\pi(\phi(\pi_{ico}(0)), \pi_{ico}(0))\right] \pi_{ico}'(0) \\ &\quad - w_x(\phi(\pi_o(0)), \pi_o(0))\, \phi'(\pi_o(0)) \left[\pi_c - \pi_o + \pi_o'(0)\right] \\ &= w_x(\phi(\pi_{ico}(0)), \pi_{ico}(0))\, \phi'(\pi_{ico}(0)) \left[\pi_{ico}'(0) - (\pi_c - \pi_o) - \pi_o'(0)\right]\end{aligned}$$

where the first inequality follows from $w(\phi(\pi_o(0)), \pi_c) - w(\phi(\pi_o(0)), \pi_o(0)) < 0$ and the last equality follows from $\pi_{ico}(0) = \pi_o(0)$ which implies that $w_\pi(\phi(\pi_{ico}(0)), \pi_{ico}(0)) = 0$. Since $w_x < 0$ and $\phi' > 0$ then it is sufficient to show that the term in square brackets is negative. To this end, we next show that $\pi_{ico}'(0) = -\infty$ and $-(\pi_c - \pi_o) - \pi_o'(0)$ is bounded.

Let's start with proving that $\lim_{\rho\to 0} \pi_{ico}'(\rho) = -\infty$. Recall that $\pi_{ico}(\rho)$ is implicitly defined by the incentive compatibility constraint

$$w(\phi(\pi_{ico}(\rho)), \pi_{ico}(\rho)) + \beta V_0(\rho) = w(\phi(\pi_{ico}(\rho)), \pi^*(\phi(\pi_{ico}(\rho)))) + \beta V_0(0).$$

Totally differentiating and evaluating at $\rho = 0$ we have that

$$w_\pi(\phi(\pi_{ico}(0)), \pi_{ico}(0))\, \pi_{ico}'(0) = -\beta V_0'(0)$$

where we used that $\pi_{ico}(0) = \pi^*(\phi(\pi_o(0)))$. Since

$$w_\pi(\phi(\pi_{ico}(0)), \pi_{ico}(0)) = w_\pi(\phi(\pi_{ico}(0)), \pi_o(0)) = 0$$

and $-\beta W_0'(0) < 0$, it must be that $\lim_{\rho\to 0} \pi_{ico}'(\rho) = -\infty$ since $w_\pi \geqslant 0$ in the relevant range.

Since $-(\pi_c - \pi_o) \leqslant \bar{\pi} - \underline{\pi}$, we are left to show that $\pi_o'(0)$ is bounded. Recall that $\pi_o(\rho)$

is the solution to

$$w_\pi\left(\phi\left(\rho\pi_c+(1-\rho)\,\pi_o\right),\pi_o\right)=0$$

and so applying the implicit function theorem we have

$$\begin{aligned}\pi_o'(0)&=-\frac{w_{\pi x}\left(\phi\left(\pi_o(0)\right),\pi_o(0)\right)x'(0)}{w_{\pi\pi}\left(\phi\left(\pi_o(0)\right),\pi_o(0)\right)}\\&=-\frac{w_{\pi x}\left(\phi\left(\pi_o(0)\right),\pi_o(0)\right)}{w_{\pi\pi}\left(\phi\left(\pi_o(0)\right),\pi_o(0)\right)}\times\frac{\phi'\left[\pi_c-\pi_o\right]}{\left[1-\phi'\pi_x^*(x)\right]}.\end{aligned}$$

Using (21) and (23) we can rewrite the expression above as

$$\pi_o'(0)=\pi_x^*(x(0))\,x'(0)=\pi_x^*(x(0))\,\frac{\phi'\left[\pi_c-\pi_o\right]}{\left[1-\phi'\pi_x^*(x)\right]}$$

so

$$|\pi_o'(0)|\leqslant|\pi_x^*(x(0))|\frac{|\phi'(\pi_o(0))|\left[\bar{\pi}-\underline{\pi}\right]}{\left[1-\phi'(\pi_o(0))\,\pi_x^*(\phi(\pi_o(0)))\right]}<\infty$$

since ϕ' is assumed to be bounded and $1-\phi'\pi_x^*(x)$ is generically not equal to zero. In particular, if the economy satisfies Assumption 2 then $\left[1-\phi'(\pi_o(0))\,\pi_x^*(\phi(\pi_o(0)))\right]>0$. Suppose not. Then it must be that $\left[1-\phi'(\pi_o(0))\,\pi_x^*(\phi(\pi_o(0)))\right]\leqslant 0$ which contradicts Assumption 2 part 3. Thus $\pi_o'(0)$ is bounded.

The above claims imply that $\left[\pi_{ico}'(0)-(\pi_c-\pi_o)-\pi_o'(0)\right]<0$ so $\Delta\omega'(0)>0$ as wanted. Q.E.D.

A.3 Assumption 4 satisfied in our examples

Lemma 3. *In our two examples, the gains from going to best response are decreasing in* x. *In general, this is true if* $\frac{w_x(\phi(\pi),\pi^*(\phi(\pi)))}{w_x(\phi(\pi),\pi)}$ *is close enough to one.*

Proof. In the Barro-Gordon model:

$$\begin{aligned}G(x)&=-\frac{1}{2}\left[\left(\psi+x-\frac{\psi+x}{2}\right)^2+\left(\frac{\psi+x}{2}\right)^2\right]+\frac{1}{2}\left[\psi^2+x^2\right]\\&=-\left(\frac{\psi+x}{2}\right)^2+\frac{1}{2}\left[\psi^2+x^2\right]\\&=-\frac{1}{4}\left[\psi^2+x^2+2\psi x\right]+\frac{1}{2}\left[\psi^2+x^2\right]\\&=\frac{1}{4}\left[\psi^2+x^2\right]-\frac{1}{2}\psi x\end{aligned}$$

so

$$G'(x)=\frac{1}{2}(x-\psi)=-\frac{1}{2}(\psi-x)$$

which is negative for all $x \in [0, \psi]$ i.e. between the Ramsey and the Markov outcome (which is the relevant range).

For the bailout example, it is more convenient to consider

$$\tilde{G}(\pi) \equiv G(\phi(\pi)).$$

Since ϕ is strictly decreasing, we have that $\tilde{G}'(\pi) = G'(\phi(\pi))\phi'(\pi)$ or $G'(x) = \tilde{G}'\left(\phi^{-1}(x)\right)/\phi'(\pi)$ so $G' \leqslant 0$ if $\tilde{G} \leqslant 0$. Note that

$$\tilde{G}(\pi) = -(1-p(\phi(\pi)))(1-\pi^*(\phi(\pi)))\psi - c(\pi^*(\phi(\pi))) + (1-p(\phi(\pi)))(1-\pi)\psi + c(\pi).$$

Since $(1-p(\phi(\pi)))\psi = c'(\pi^*(\phi(\pi)))$ we can write

$$\begin{aligned}\tilde{G}'(\pi) &= p'(\phi(\pi))\phi'(\pi)(1-\pi^*(\phi(\pi)))\psi - p'(\phi(\pi))\phi'(\pi)(1-\pi)\psi - (1-p(\phi(\pi)))\psi + c'(\pi) \\ &= p'(\phi(\pi))\phi'(\pi)(\pi-\pi^*(\phi(\pi)))\psi - \left[(1-p(\phi(\pi)))\psi - c'(\pi)\right].\end{aligned}$$

Recall that

$$\pi^*(e) = \frac{(1-p(e))\psi}{\lambda}.$$

Thus,

$$\begin{aligned}\tilde{G}(\pi) &= p'(\phi(\pi))\phi'(\pi)\left(\pi - \frac{(1-p(\phi(\pi)))\psi}{\lambda}\right)\psi - [(1-p(\phi(\pi)))\psi - \lambda\pi] \\ &= \left[\frac{(1-p(\phi(\pi)))\psi}{\lambda} - \pi\right]\left[-p'(\phi(\pi))\phi'(\pi) - \frac{\lambda}{\psi}\right]\psi.\end{aligned}$$

We are now going to show that the first term in square brackets is positive while the second is negative. Let's start with the first. Since we considering $\pi \in [0, \pi_o(0)]$ it must be that for any ψ, λ, and interior π, the first term is positive since $\pi < \pi^*(\phi(\pi))$. Consider next the second term, $[-p'(\phi(\pi))\phi'(\pi) - \lambda/\psi]$. Note that $-p'(\phi(\pi))\phi'(\pi) > 0$ and it is increasing in π. Moreover, $\pi_M(\psi)$ is increasing in ψ. These two observations imply that we can find a ψ sufficiently small such that the second term is negative for all $\pi \in [0, \pi_M(\psi)]$. Thus, for ψ small enough, we have that $\tilde{G}'(\pi) \leqslant 0$.

In general, we have that

$$\tilde{G}'(\pi) = [w_x(\phi(\pi), \pi^*(\phi(\pi))) - w_x(\phi(\pi), \pi)]\phi'(\pi) - w_\pi(\phi(\pi), \pi)$$

which is negative if

$$w_x(\phi(\pi), \pi)\left[\frac{w_x(\phi(\pi), \pi^*(\phi(\pi)))}{w_x(\phi(\pi), \pi)} - 1\right]\phi'(\pi) - w_\pi(\phi(\pi), \pi) \leqslant 0$$

which is true if $\frac{w_x(\phi(\pi),\pi^*(\phi(\pi)))}{w_x(\phi(\pi),\pi)}$ is close enough to one. Q.E.D.

A.4 Proof of Proposition 3

The problem for the rule designer for a generic horizon $k+1$ can be written as

$$W_{k+1}(\rho) = \max_{\pi_c, x, \sigma \in [0,1]} (\rho + (1-\rho)\sigma)\left[w(x,\pi_c) + \beta W_k\left(\frac{\rho}{\rho + (1-\rho)\sigma}\right)\right] \\ + (1-\rho)(1-\sigma)\left[w(x,\pi^*(x)) + \beta W_k(0)\right] \tag{36}$$

subject to the implementability condition,

$$x = \phi\left((\rho + (1-\rho)\sigma)\pi_c + (1-\rho)(1-\sigma)\pi^*(x)\right),$$

and the incentive compatibility constraint for the optimizing type,

$$\sigma\left[w(x,\pi_c) + \beta_o V_k\left(\frac{\rho}{\rho + (1-\rho)\sigma}\right) - \left[w(x,\pi^*(x)) + \beta_o V_k(0)\right]\right] = 0.$$

The value for the optimizing type for a generic horizon $k+1$ is

$$V_{k+1}(\rho) = \sigma_k(\rho)\left[w(x_{k+1}(\rho), \pi_{c,k+1}(\rho)) + \beta_o V_k\left(\frac{\rho}{\rho + (1-\rho)\sigma}\right)\right] \\ + (1-\sigma_k(\rho))\left[w(x_{k+1}(\rho), \pi^*(x_{k+1}(\rho))) + \beta_o V_k(0)\right].$$

With this setup, we can now turn to the proof of the proposition (as usual we consider the case with $w_x < 0$):

Part 1. For $\rho = 0$ it is clear that the equilibrium is the repetition of the static outcome i.e. all k, $W_k(0) = W_0(0)/(1-\beta)$ and $V_k(0) = V_0(0)/(1-\beta_o)$, since no incentives can be provided to the optimizing type.

Part 3. Consider $\rho \in (\rho_1^*, 1]$. Here we know that in the twice repeated problem it is optimal to separate in the first period. We now show it is also optimal to separate for all horizons $k \geqslant 2$. Consider any horizon $k+1$ with $k \geqslant 1$. Suppose it is optimal to separate for horizons $0, 1, ..., k$ at prior ρ. We next show it is optimal to separate in $k+1$. Define

$$\Delta V_k(\rho) \equiv V_k(\rho) - V_k(0).$$

Note that regardless of the horizon, if there is separation next period we have that

$$\begin{aligned}\Delta V_k(\rho) &= \left[V_0(\rho) + \left(\beta + \beta^2 + ... + \beta^k\right) V_0(0)\right] - \left(1 + \beta + \beta^2 + ... + \beta^k\right) V_0(0) \\ &= V_0(\rho) - V_0(0)\end{aligned}$$

Thus, by the same argument used in Lemma 7 for the case $k = 1$, pooling with probability one is preferable to pooling with some probability $\sigma < 1$. Consequently the IC constraint for the optimizing type is the same as in $k = 1$ case and so $x_{ico,k+1}(\rho) = x_{ico,1}(\rho)$. Thus, the only relevant options are pooling or separating with probability one. We can just compare two values:

$$\begin{aligned}W_{k+1}^{pool}(\rho) &= w\left(x_{ico,1}(\rho), \phi^{-1}(x_{ico,1}(\rho))\right) + \beta W_0(\rho) \\ &\quad + \left(\beta^2 + \beta^3 + ... + \beta^{k+1}\right)[\rho W_0(1) + (1-\rho) W_0(0)], \\ W_{k+1}^{sep}(\rho) &= W_0(\rho) + \left(\beta + \beta^2 + \beta^3 + ... + \beta^{k+1}\right)[\rho W_0(1) + (1-\rho) W_0(0)]\end{aligned}$$

where we used that $x_{ico,k+1}(\rho) = x_{ico,1}(\rho)$. Therefore,

$$\begin{aligned}W_{k+1}^{sep}(\rho) - W_{k+1}^{pool}(\rho) &= \left[W_0(\rho) - w\left(x_{ico,1}(\rho), \phi^{-1}(x_{ico,1}(\rho))\right)\right] \\ &\quad - \beta\{W_0(\rho) - [\rho W_0(1) + (1-\rho) W_0(0)]\} \\ &= W_1^{sep}(\rho) - W_1^{pool}(\rho) > 0\end{aligned}$$

Thus, for $\rho > \rho_1^*$ it is always optimal to separate for any horizon.

Part 2. We now turn to show that for reputation low enough it is optimal to pool with no randomization. Note that if it is optimal to pool – at least partially – then $x_k = x_{ico,k}(\rho)$ and it solves

$$G(x_{ico,k}(\rho)) = \beta \Delta V_{k-1}\left(\frac{\rho}{\rho + (1-\rho)\sigma}\right) \tag{37}$$

where recall $\Delta V_k(\rho) \equiv V_k(\rho) - V_k(0)$ and $G(x) = w(x, \pi^*(x)) - w\left(x, \phi^{-1}(x)\right)$.

First, note that there exists $\hat{\rho} < \rho_1^*$ such that $\Delta V_1(\rho) \geqslant \Delta V_0(\rho)$ for all $\rho \in [0, \hat{\rho}]$ as showed in Figure 5. To see this, note that

$$\Delta V_0(\rho) = \Delta V^{sep}(\rho) = w(x_0(\rho), \pi^*(x_0(\rho))) - w(x_0(0), \pi^*(x_0(0)))$$

and, using the result in Proposition 2 that for $\rho \leqslant \rho_1^*$ it is optimal to pool without random-

Figure 5: Dynamic incentives, $\Delta V_k(\rho)$

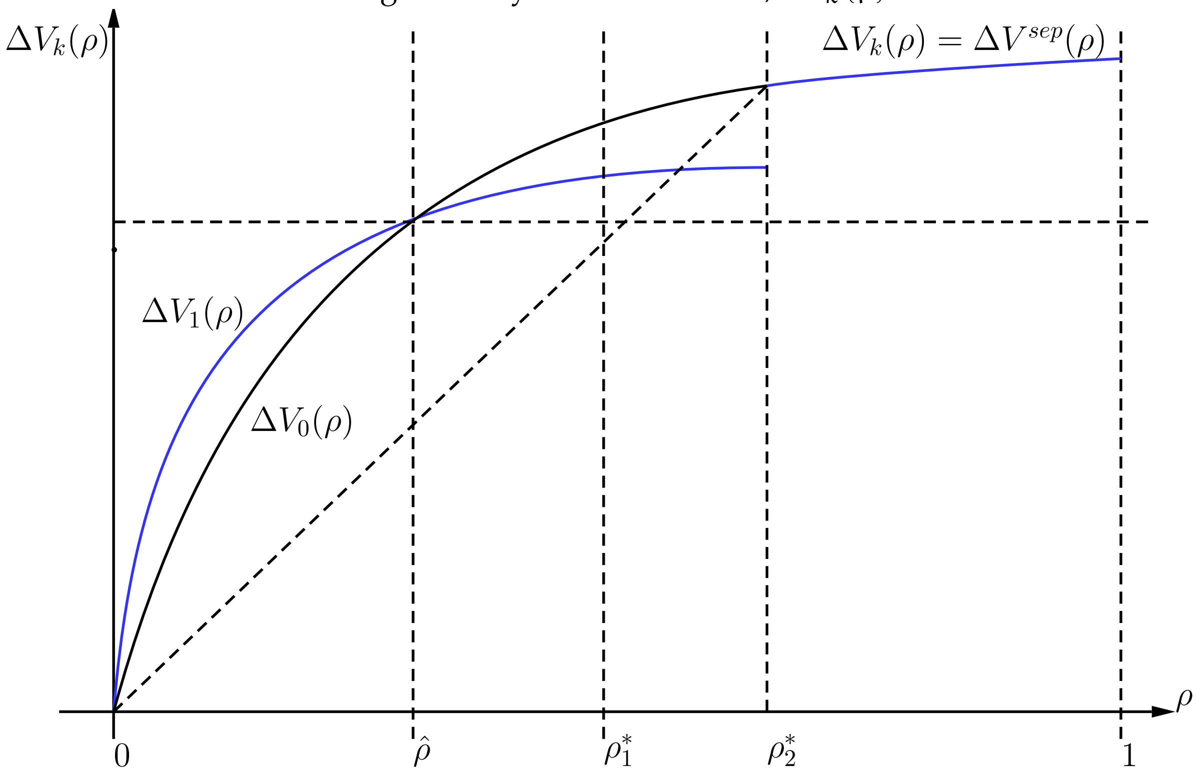

ization,

$$\begin{aligned}\Delta V_1(\rho) &= w\left(x_{ico,1}(\rho), \phi^{-1}(x_{ico,1}(\rho))\right) + \beta_o V_0(\rho) - (1+\beta_o) V_0(0) \\ &= w(x_{ico,1}(\rho), \pi^*(x_{ico,1}(\rho))) + \beta_o V_0(0) - (1+\beta_o) V_0(0) \\ &= w(x_{ico,1}(\rho), \pi^*(x_{ico,1}(\rho))) - w(x_0(0), \pi^*(x_0(0)))\end{aligned}$$

where the first equality is the definition of ΔV_1, the second uses the binding incentive constraint, and the last line is just algebra. Comparing the two expressions above, it follows that $\Delta V_1(\rho) > \Delta V_0(\rho)$ if and only if $w(x_{ico,1}(\rho), \pi^*(x_{ico,1}(\rho))) > w(x_0(\rho), \pi^*(x_0(\rho)))$ or equivalently

$$x_{ico,1}(\rho) < x_0(\rho).$$

From the proof of Proposition 2, we know that

$$x_{ico,1}(0) = x_0(0)$$

and for some finite $M > 0$,

$$-\infty = \lim_{\rho\to 0} \frac{\partial x_{ico,1}(\rho)}{\partial \rho} < -M < \lim_{\rho \to 0} \frac{\partial x_0(\rho)}{\partial \rho}.$$

Thus, by continuity, there must exists $\hat{\rho} \leqslant \rho_1^*$ such that $x_{ico,1}(\rho) < x_0(\rho)$ and therefore $\Delta V_1(\rho) \leqslant \Delta V_0(\rho)$ for all $\rho \leqslant \hat{\rho}$.

Next, we show that for a generic horizon $k+1$, if $\Delta V_k(\rho) \geqslant \Delta V_{k-1}(\rho) \geqslant \Delta V_0(\rho)$, $x_{ico,k}(\rho) \leqslant x_{ico,k-1}(\rho) \leqslant x_{ico,1}(\rho)$, and it is optimal to pool for $k, k-1, ..., 1$ for all $\rho \leqslant \hat{\rho}$ then for $k+1$: i) it is not optimal to randomize, ii) it is optimal to pool, iii) $\Delta V_{k+1}(\rho) \geqslant \Delta V_k(\rho)$ and $x_{ico,k+1}(\rho) \leqslant x_{ico,k}(\rho)$.

To see that randomization is not optimal, we just have to follow the same steps in Lemma 7, noting that if it is optimal to randomize then it is optimal to choose σ sufficiently low so that $\rho'(\pi_c|\pi_c,\sigma) \geqslant \rho_2^*$ to take advantage of the discontinuous jump in ΔV_k at ρ_2^* and relax the incentive constraint. Using this, we have

$$\begin{aligned}
&[\rho+(1-\rho)\sigma]\,\Delta V_k\left(\frac{\rho}{\rho+(1-\rho)\sigma}\right)+(1-\rho)(1-\sigma)\,\Delta V_k(0) \qquad (38)\\
&=[\rho+(1-\rho)\sigma]\left[V_0\left(\frac{\rho}{\rho+(1-\rho)\sigma}\right)-V_0(0)\right]+(1-\rho)(1-\sigma)\,[V_0(0)-V_0(\rho)]\\
&\leqslant[\rho+(1-\rho)\sigma]\,[V_0(\rho)-V_0(0)]\\
&\leqslant[\rho+(1-\rho)\sigma]\,[V_k(\rho)-V_k(0)]
\end{aligned}$$

where the first equality follows from the fact that $\rho/[\rho+(1-\rho)\sigma] \geqslant \rho_2^*$ and the optimality of separation with probability one for $\rho \geqslant \rho_2^*$, the second inequality follows from concavity of V_0, and the last one from $\Delta V_k(\rho) \geqslant \Delta V_0(\rho)$ for $\rho \in [0,\hat{\rho}]$. Using (38) in the expression (57) in Lemma 7 gives that randomization is not optimal.

Thus the only relevant options are pooling or separating with probability one. We can just compare two values:

$$W_{k+1}(\rho) = \max\left\{W_{k+1}^{pool}(\rho), W_{k+1}^{sep}(\rho)\right\} \qquad (39)$$

where

$$W_{k+1}^{sep}(\rho) = W_0(\rho) + \beta\,[\rho W_k(1) + (1-\rho)\,W_k(0)] \qquad (40)$$

$$W_{k+1}^{pool}(\rho) = w\left(x_{ico,k+1}(\rho), \phi^{-1}(x_{ico,k}(\rho))\right) + \beta W_k(\rho) \qquad (41)$$

where

$$G(x_{ico,k+1}(\rho)) = \beta\Delta V_k(\rho) \qquad (42)$$

To see that it is optimal to pool, note that by our induction hypothesis

$$G(x_{ico,k}(\rho)) = \beta\Delta V_{k-1}(\rho)$$

and $\Delta V_{k-1}(\rho) \leqslant \Delta V_k(\rho)$. Thus, since G is decreasing, it follows that

$$x_{ico,k+1}(\rho) \leqslant x_{ico,k}(\rho) \leqslant x_{ico,1}(\rho). \tag{43}$$

Thus,

$$\begin{aligned}
W_{k+1}^{pool}(\rho) - W_{k+1}^{sep}(\rho) &= w\left(x_{ico,k+1}, \phi^{-1}(x_{ico,k+1})\right) + \beta W_k^{pool}(\rho) \\
&\quad - \left\{W_0(\rho) + \frac{\beta(1-\beta^k)}{1-\beta}[\rho W_0(1) + (1-\rho)W_0(0)]\right\} \\
&\geqslant w\left(x_{ico,k+1}, \phi^{-1}(x_{ico,k+1})\right) \\
&\quad + \beta\left\{W_0(\rho) + \frac{\beta(1-\beta^{k-1})}{1-\beta}[\rho W_0(1) + (1-\rho)W_0(0)]\right\} \\
&\quad - \left\{W_0(\rho) + \frac{\beta(1-\beta^k)}{1-\beta}[\rho W_0(1) + (1-\rho)W_0(0)]\right\} \\
&= w\left(x_{ico,k+1}, \phi^{-1}(x_{ico,k+1})\right) - W_0(\rho) + \beta[W_0(\rho) - \rho W_0(1) + (1-\rho)W_0(0)] \\
&\geqslant w\left(x_{ico,1}, \phi^{-1}(x_{ico,1})\right) - W_0(\rho) + \beta[W_0(\rho) - \rho W_0(1) + (1-\rho)W_0(0)] \\
&\geqslant 0
\end{aligned}$$

where the first equality just uses the definitions of W_{k+1}^{pool} and W_{k+1}^{sep}, the second inequality follows from $W_k^{pool} \geqslant W_k^{sep}$ which follows from the induction hypothesis, the third equality follows from algebra, and the fourth inequality follows from (43), and the last one from $\rho \leqslant \hat{\rho} \leqslant \rho_1^*$ which implies that pooling is optimal in the twice repeated economy. Thus it is optimal to pool.

That $\Delta V_{k+1}(\rho) \geqslant \Delta V_k(\rho) \geqslant \Delta V_0(\rho)$ follows directly from (43).

We thus established that for $\rho \in [0, \hat{\rho}]$, for all k it is optimal to pool and $\{x_k\} = \{x_{ico,k}(\rho)\}$ is a decreasing sequence in $[\underline{\pi}, \bar{\pi}]$ thus it must converge and its limit is given by x_{CK} defined in the text. Q.E.D.

Note that if $W_0(\rho)$ is concave then uncertainty is beneficial at any horizon $k \geqslant 0$, in that $W_k(\rho) > \rho W_k(1) + (1-\rho)W_k(0)$ for all k as shown in Figure 6. This is true despite the fact that for a general horizon k, $W_k(\rho)$ is the upper envelope of two (concave) functions, $W_k(\rho) = \max\left\{W_k^{pool}(\rho), W_k^{sep}(\rho)\right\}$, and therefore it is not globally concave.

Lemma 4. *If* $W_k(\rho) > \rho W_k(1) + (1-\rho)W_k(0)$ *then* $W_{k+1}(\rho) > \rho W_{k+1}(1) + (1-\rho)W_{k+1}(0)$.

Figure 6: Equilibrium values: uncertainty is beneficial for $k \geqslant 1$

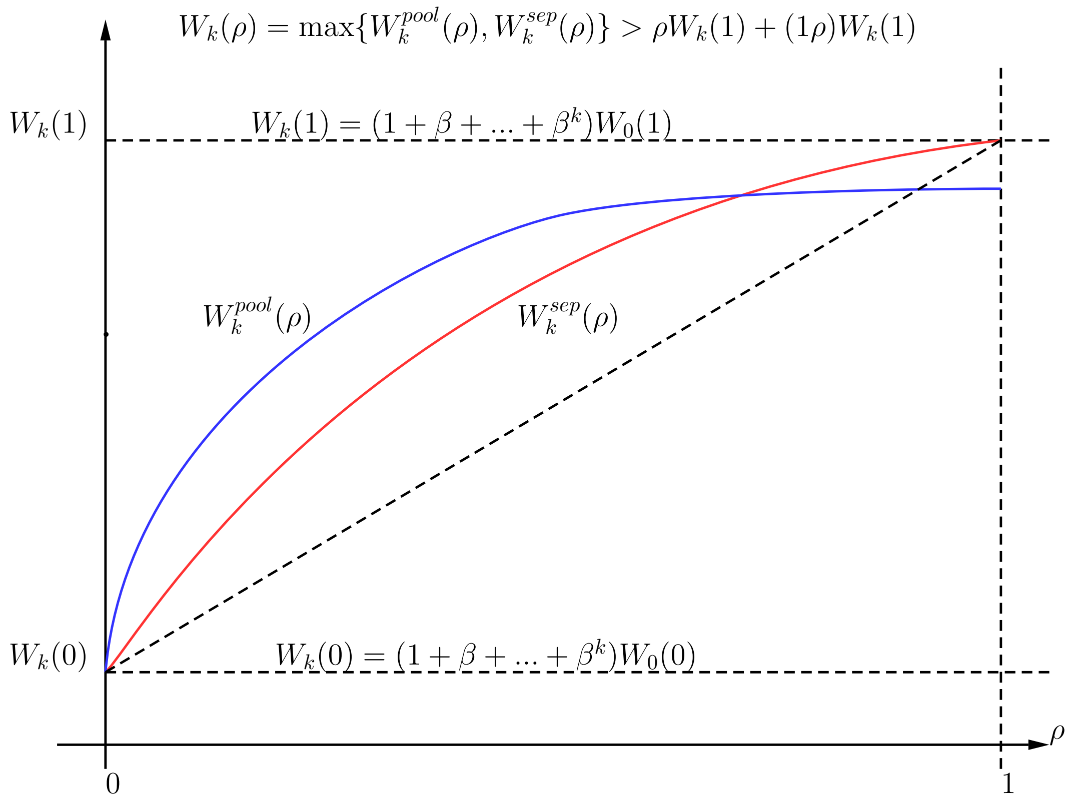

Proof. Note that

$$\begin{aligned} W_{k+1}(\rho) &\geqslant W_{k+1}^{sep}(\rho) \\ &= W_0(\rho) + \beta\rho W_k(1) + \beta(1-\rho) W_k(0) \\ &\geqslant [\rho W_0(1) + (1-\rho) W_0(0)] + \beta\rho W_k(1) + \beta(1-\rho) W_k(0) \\ &= \rho W_{k+1}(1) + (1-\rho) W_{k+1}(0) \end{aligned}$$

where the first inequality follows from the definition of W_{k+1} in (39), the second line follows from the definition of $W_{k+1}^{sep}(\rho)$ in (40), the third line follows from uncertainty being beneficial in the static problem and the induction hypothesis, and the last line follows from the fact that for $\rho \in \{0, 1\}$, $W_k(\rho) = W_{k+1}(\rho) = W_0(\rho)$. Q.E.D.

The above lemma and the fact that W_0 satisfies the property implies that for all k, $W_k(\rho)$ satisfies the above property.

A.5 Proof of Proposition 4

Consider first $\rho \in [0, \rho_1^*]$. From Proposition 2, we know that if signals are perfectly informative, it is optimal to be in the pooling regime so $\pi_o = \pi_c$. As we proved in the Lemma

7, conditional on pooling, it is preferable to choose $\sigma_\varepsilon = 0$ to relax the incentive constraint (13).

Next, we show that pooling and $\sigma_\varepsilon = 0$ is preferable to a separating outcome with $\sigma_\varepsilon > 0$. Suppose not and let $W^{sep}(\sigma_\varepsilon > 0)$ be the value attained by separating and choosing $\sigma_\varepsilon > 0$. Consider a deviation with pooling and

$$x^{dev} = \phi(\rho\pi_c + (1-\rho)\pi_o) > x_{ico},$$
$$\pi_o^{dev} = \pi_c^{dev} = \pi^{dev} = \rho\pi_c + (1-\rho)\pi_o$$

where π_c and π_o are the policies associated with separation and $\sigma_\varepsilon > 0$. Clearly, the expected continuation value of separating and choosing $\sigma_\varepsilon > 0$ is lower than the continuation value of the proposed deviation, $W_0(\rho)$, due to the concavity of W_0 and the fact that the posteriors are martingales. Lets compare the static values. By the concavity of w

$$\rho w(x, \pi_c) + (1-\rho) w(x, \pi_o) < w\left(x^{dev}, \pi^{dev}\right)$$

so the static value is higher in the proposed deviation. We are thus left to check that the deviation is incentive compatible or

$$w\left(x^{dev}, \pi^{dev}\right) + \beta_o V_0(\rho) \geqslant w\left(x^{dev}, \pi^*\right) + \beta_o V_0(\rho)$$

$$\iff \beta_o [V_0(\rho) - V_0(0)] \geqslant w\left(x^{dev}, \pi^*\right) - w\left(x^{dev}, \pi^{dev}\right) = G\left(x^{dev}\right).$$

We have

$$\begin{aligned}\beta_o [V_0(\rho) - V_0(0)] &\geqslant w(x_{ico}, \pi^*(x_{ico})) - w(x_{ico}, \pi_{ico}) \\ &= G(x_{ico}) > G\left(x^{dev}\right) = w\left(x^{dev}, \pi^*\right) - w\left(x^{dev}, \pi^{dev}\right)\end{aligned}$$

since $x^{dev} \geqslant x_{ico}$ and G is decreasing by Assumption 4. Thus the deviation is incentive compatible and it attains a higher value than $W^{sep}(\sigma_\varepsilon > 0)$. Therefore, it is optimal to pool and set $\sigma_\varepsilon = 0$ for $\rho \in [0, \rho_1^*]$.

The argument for $\rho \in [\rho_2^*, 1]$ is provided in the main text. Q.E.D.

A.6 Proof of Proposition 5

We start by proving that for ρ close to 1 it is optimal to have stochastic rules. Consider a ρ close enough to 1 so that with a deterministic rule it is optimal to separate so $\pi = \pi_0(\rho)$

. The value of this policy is

$$\begin{aligned}W &= [\rho w(x_0, \pi_0) + (1-\rho) w(x_0, \pi^*(x_0))] + \beta [\rho W_0(1) + (1-\rho) W_0(0)] \\ &= W_0(\rho) + \beta[\rho W_0(1) + (1-\rho) W_0(0)]\end{aligned}$$

We now show that if ρ is close to 1 then this policy is dominated by one that calls for the commitment type to play the static best response with some positive probability. Consider a deviation indexed by $\varepsilon > 0$ sufficiently small so that

$$\pi_c = \begin{cases} \pi_0(\rho) & \text{with pr } 1-\varepsilon \\ 1 & \text{with pr } \varepsilon \end{cases}$$

so after observing a bailout the posterior is

$$\rho' = \frac{\rho\varepsilon}{\rho\varepsilon + (1-\rho)} = \frac{\rho}{\rho + (1-\rho)/\varepsilon} > 0$$

and after no-bailout $\rho' = 1$. The value of this deviation is then

$$\begin{aligned}W^{\text{dev}}(\varepsilon) = [\rho(1-\varepsilon) w(x_0(\varepsilon), \pi_0) + [\rho\varepsilon + (1-\rho)] w(x_0(\varepsilon), \pi^*(x_0))] \\ + \beta\left[\rho(1-\varepsilon) W_0(1) + [\rho\varepsilon + (1-\rho)] W_0(\rho')\right]\end{aligned}$$

Since $W = W^{\text{dev}}(0)$ we have

$$\begin{aligned}W^{\text{dev}}(\varepsilon) - W^{\text{dev}}(0) &= \Delta\omega(\varepsilon) + \beta\Delta\Omega(\varepsilon) \\ &\approx \left[\Delta\omega'(\varepsilon) + \beta\Delta\Omega'(\varepsilon)\right]\varepsilon\end{aligned}$$

where

$$\Delta\omega(\varepsilon) = [\rho(1-\varepsilon) w(x_0(\varepsilon), \pi_0) + [\rho\varepsilon + (1-\rho)] w(x_0(\varepsilon), \pi^*(x_0))]$$

$$\Delta\Omega(\varepsilon) = \left[\rho(1-\varepsilon) W_0(1) + [\rho\varepsilon + (1-\rho)] W_0(\rho')\right] - [\rho W_0(1) + (1-\rho) W_0(0)]$$

Note that

$$\begin{aligned}\Delta\Omega'(\varepsilon) &= -\rho W_0(1) + \rho W_0(\rho'(\varepsilon)) + [\rho\varepsilon + (1-\rho)] W_0'(\rho'(\varepsilon)) \frac{\partial\rho'}{\partial\varepsilon} \\ &= -\rho\left[W_0(1) - W_0(\rho'(\varepsilon))\right] + [\rho\varepsilon + (1-\rho)] W_0'(\rho'(\varepsilon)) \frac{\partial\rho'}{\partial\varepsilon}\end{aligned}$$

As $\varepsilon \to 0$

$$\Delta\Omega'(\varepsilon) \to -\rho[W_0(1) - W_0(0)] + [(1-\rho)] W_0'(0) \frac{\partial\rho'}{\partial\varepsilon}$$

$$\frac{\partial \rho'}{\partial \varepsilon} = \frac{\rho\left[\rho\varepsilon + (1-\rho)\right] - \rho\varepsilon\rho}{\left[\rho\varepsilon + (1-\rho)\right]^2} \to \frac{\rho\left(1-\rho\right)}{\left(1-\rho\right)^2} = \frac{\rho}{\left(1-\rho\right)}$$

so for ρ close to one

$$\lim_{\rho\to 1}\lim_{\varepsilon\to 0}\Delta\Omega'\left(\varepsilon\right) = \infty$$

Thus to show that the deviation is profitable it is sufficient to show that $\Delta\omega > -M$ for some M sufficiently large. Consider

$$\begin{aligned}\Delta\omega'\left(\varepsilon\right) &= \rho\left[w\left(x_0\left(\varepsilon\right), \pi^*\left(x_0\left(\varepsilon\right)\right)\right) - w\left(x_0\left(\varepsilon\right), \pi_0\right)\right] \\ &\quad + \left\{\rho\left(1-\varepsilon\right) w_x\left(x_0\left(\varepsilon\right), \pi_0\right) + \left[\rho\varepsilon + (1-\rho)\right] w_x\left(x_0\left(\varepsilon\right), \pi^*\left(x_0\right)\right)\right\} \frac{\partial x_0\left(\varepsilon\right)}{\partial \varepsilon} \\ &\quad + \rho\left(1-\varepsilon\right) w_\pi\left(x_0\left(\varepsilon\right), \pi_0\right)\end{aligned}$$

where

$$\frac{\partial x_0\left(\varepsilon\right)}{\partial \varepsilon} = \phi_\pi \rho\left(\pi^* - \pi_0\right)$$

Since the first term in square brackets is positive we have that

$$\begin{aligned}\Delta\omega'\left(\varepsilon\right) &> \left\{\rho\left(1-\varepsilon\right) w_x\left(x_0\left(\varepsilon\right), \pi_0\right) + \left[\rho\varepsilon + (1-\rho)\right] w_x\left(x_0\left(\varepsilon\right), \pi^*\left(x_0\right)\right)\right\} \frac{\partial x_0\left(\varepsilon\right)}{\partial \varepsilon} \\ &\quad + \rho\left(1-\varepsilon\right) w_\pi\left(x_0\left(\varepsilon\right), \pi_0\right) \\ &\quad = \left\{\rho w_x\left(x_0, \pi_0\right) + \left(1-\rho\right) w_x\left(x_0, \pi^*\left(x_0\right)\right)\right\} \phi_\pi \rho\left(\pi^* - \pi_0\right) + \rho w_\pi\left(x_0, \pi_0\right)\end{aligned}$$

with w_x and ϕ_π bounded, as $\rho \to 1$ we have that the last expression converges to

$$w_x\left(x_0, \pi_0\right)\phi_\pi\left(\pi^* - \pi_0\right) + w_\pi\left(x_0, \pi_0\right) > w_x\left(x_0, \pi_0\right)\phi_\pi\left(\pi^* - \pi_0\right) > w_x\left(x_0, \pi_0\right)\phi_\pi\left(\bar{\pi} - \underline{\pi}\right) > -M$$

for some $M < \infty$.

(Notice that to derive this result we are not relying on the concavity of W_0 but only: i) $W_0' > 0$ and ii) properties of Bayes' rule.)

We now prove that for ρ close to zero a deterministic rule is optimal. Since W_0 is concave and the posterior is a martingale,

$$\int \rho'\left(\pi, \rho\right)\left[\rho\sigma_c\left(\pi\right) + \left(1-\rho\right)\sigma_o\left(\pi\right)\right] d\pi = \rho,$$

then

$$W_0\left(\rho\right) \geqslant \int W_0\left(\rho'\left(\pi, \rho\right)\right)\left[\rho\sigma_c\left(\pi\right) + \left(1-\rho\right)\sigma_o\left(\pi\right)\right] d\pi.$$

Thus randomization can be optimal only if it improves that static outcome by reducing x.

Then it must be that

$$w\left(\phi\left(\pi_{ico}\right),\pi_{ico}\right) < \int w\left(x,\pi\right)\left[\rho\sigma_c\left(\pi\right) + \left(1-\rho\right)\sigma_o\left(\pi\right)\right] d\pi$$

where x is given by (16). A necessary condition is that

$$x < \phi\left(\pi_{ico}\right) \iff \mathbb{E}\pi \equiv \int \pi\left[\rho\sigma_c\left(\pi\right) + \left(1-\rho\right)\sigma_o\left(\pi\right)\right] d\pi < \pi_{ico}$$

Thus, it is sufficient to show that $\mathbb{E}\pi > \pi_{ico}$ to prove our result. Suppose by way of contradiction that it is not optimal to have π_{ico} with probability 1 so

$$\int \pi\left[\rho\sigma_c\left(\pi\right) + \left(1-\rho\right)\sigma_o\left(\pi\right)\right] d\pi < \pi_{ico} \tag{44}$$

and since we consider $\rho \to 0$ then

$$\mathbb{E}_o\pi = \int \pi\sigma_o\left(\pi\right) d\pi \leqslant \pi_{ico} \tag{45}$$

otherwise we can make ρ arbitrary close to 0 so that the inequality in (44) is reversed. From the incentive constraint, it must be that $\forall\pi \in \operatorname{Supp}\sigma_0$

$$w\left(x,\pi\right) + \beta V_0\left(\rho'\left(\pi,\rho\right)\right) \geqslant w\left(x,\pi^*\left(x\right)\right) + \beta V_0\left(\rho'\left(\pi^*\left(x\right),\rho\right)\right) \geqslant w\left(x,\pi^*\left(x\right)\right) + \beta V_0\left(0\right) \tag{46}$$

where the second inequality follows from V_0 being increasing in the posterior and $\rho'\left(\pi^*\left(x\right),\rho\right) \geqslant 0$. Note now that by properties of Bayes' rule

$$\int \rho'\left(\pi,\rho\right)\sigma_o\left(\pi\right) d\pi = \int \frac{\rho\sigma_c\left(\pi\right)}{\rho\sigma_c\left(\pi\right) + \left(1-\rho\right)\sigma_o\left(\pi\right)}\sigma_o\left(\pi\right) d\pi \leqslant \rho. \tag{47}$$

Thus we have:

$$\begin{aligned}\int \left[w\left(x,\pi\right) + \beta V_0\left(\rho'\left(\pi,\rho\right)\right)\right]\sigma_o\left(\pi\right) d\pi &< w\left(x,\mathbb{E}_o\pi\right) + \beta V_0\left(E_o\rho'\right) \\ &\leqslant w\left(x,\mathbb{E}_o\pi\right) + \beta V_0\left(\rho\right) \\ &= w\left(\phi\left(\mathbb{E}_o\pi\right),\mathbb{E}_o\pi\right) + \beta V_0\left(\rho\right)\end{aligned} \tag{48}$$

where the first inequality follows from the strict concavity of w (in π) and V_0, the second inequality from (47) and V_0 strictly increasing. Thus combining (46) and (48) we have that

$$w\left(\phi\left(\mathbb{E}_o\pi\right),\mathbb{E}_o\pi\right) + \beta V_0\left(\rho\right) > w\left(\phi\left(\mathbb{E}_o\pi\right),\pi^*\left(\phi\left(\mathbb{E}_o\pi\right)\right)\right) + \beta V_0\left(0\right)$$

Since π_{ico} is the smallest solution to

$$w(\phi(\pi), \pi) + \beta V_0(\rho) = w(\phi(\pi), \pi^*(\pi)) + \beta V_0(0)$$

then it follows that for ρ close to zero

$$\pi_{ico}(\rho) < \mathbb{E}_o \pi$$

a contradiction. Q.E.D.

A.7 Proof of Proposition 7

Suppose by way of contradiction that pooling is optimal so $\pi_o = \pi_c$. Consider a deviation that leaves x, π_o, and $V_o(\pi_o)$ constant and set $\pi_c = \arg\max_{\pi_c} w(x, \pi_c)$ and $V_o(\pi_c) = \arg\max_V \bar{W}_0(V, 1)$. Note that this deviation is feasible and attains a higher current value. Next we show that the expected continuation value is also higher. Under the conjectured optimum, the rule designer's expected continuation value is

$$\bar{W}_0(V_o, \rho) = \max_{x, \pi_c, \pi_o} \rho w(x, \pi_c) + (1-\rho) w(x, \pi_o)$$

subject to

$$w(x, \pi_o) \geqslant w(x, \pi) \ \forall \pi,$$
$$w(x, \pi_o) = V_o.$$

Under our proposed deviation, the rule designer's expected continuation value is

$$\rho W_0(1) + (1-\rho) \bar{W}_0(V_c, 0) = \max_{x_c, x_o, \pi_c, \pi_o} \rho w(x_c, \pi_c) + (1-\rho) w(x_o, \pi_o)$$

subject to

$$w(x, \pi_o) \geqslant w(x, \pi) \ \forall \pi,$$
$$w(x, \pi_o) = V_o.$$

Therefore

$$\bar{W}_0(V_c, \rho) \leqslant \rho W_0(1) + (1-\rho) \bar{W}_0(V_c, 0)$$

since the set of feasible outcomes in the first problem is a subset of the second because the former has the implicit additional constraint that $x_c = x_o = x$. Q.E.D.

B Example where uncertainty is not beneficial

FDI example Consider an economy populated by workers and foreign investors. Workers inelastically supply one unit of labor and have preferences over consumption given by $v(c)$. Investors are risk neutral and have a large endowment e that can be either consumed or invested. Let k be the amount invested. Output is produced by competitive firms with a production function $y = k^{\alpha} l^{1-\alpha}$ for $\alpha \in (0,1)$. The government cares about the welfare of the workers and can tax the investors' capital income and lump-sum rebate the proceeds to the workers. Let π be the tax rate. The government's preferences are

$$w(k,\pi) = \omega(k) + \pi R(k) k$$

where $\omega(k) = (1-\alpha) k^{\alpha}$ and $R(k) = \alpha k^{\alpha-1}$ are the competitive factor prices. The problem for an individual investor i is

$$\max_{k_i} u(k_i, k, \pi) = \max_{k_i} e - k_i + (1-\pi) R(k) k_i$$

Optimality and representativeness imply that in equilibrium

$$k(\pi) = \phi(\pi) = \alpha^{1/(1-\alpha)} (1-\pi)^{1/(1-\alpha)}. \tag{49}$$

It is easy to show that $\pi^*(k) = 1$ for all k and in all computed examples $\pi_c(\rho) = 0$ for all ρ.

From (49), it follows that $k(\pi)$ is convex in π as

$$k''(\pi) = \alpha^{1/(1-\alpha)} \frac{\alpha}{1-\alpha} \frac{1}{1-\alpha} (1-\pi)^{\frac{\alpha}{1-\alpha}-1} > 0$$

as shown in the first panel of Figure 7. That is, when (expected) taxes are high a reduction in taxes results in a smaller increase in investment than when taxes are low. This is because the equilibrium interest rate is more sensitive when taxes are high and thus the effect of the tax reduction on investment is mitigated by the steep reduction in the interest rate. Thus, the complementarity between taxes and investment is lower when taxes are high (far away from the Ramsey policy).

The failure of having a concave ϕ is not enough to rule out the concavity of $W_0(\rho)$ on its own. As we show in the fourth panel of Figure 7, if α is large enough so that preferences $w(k,\pi)$ are close enough to be linear then $W_0(\rho)$ is convex for low levels of reputation but as we increase the concavity of w by reducing α then $W_0(\rho)$ is globally concave.

Figure 7: Static value and private action for the FDI example

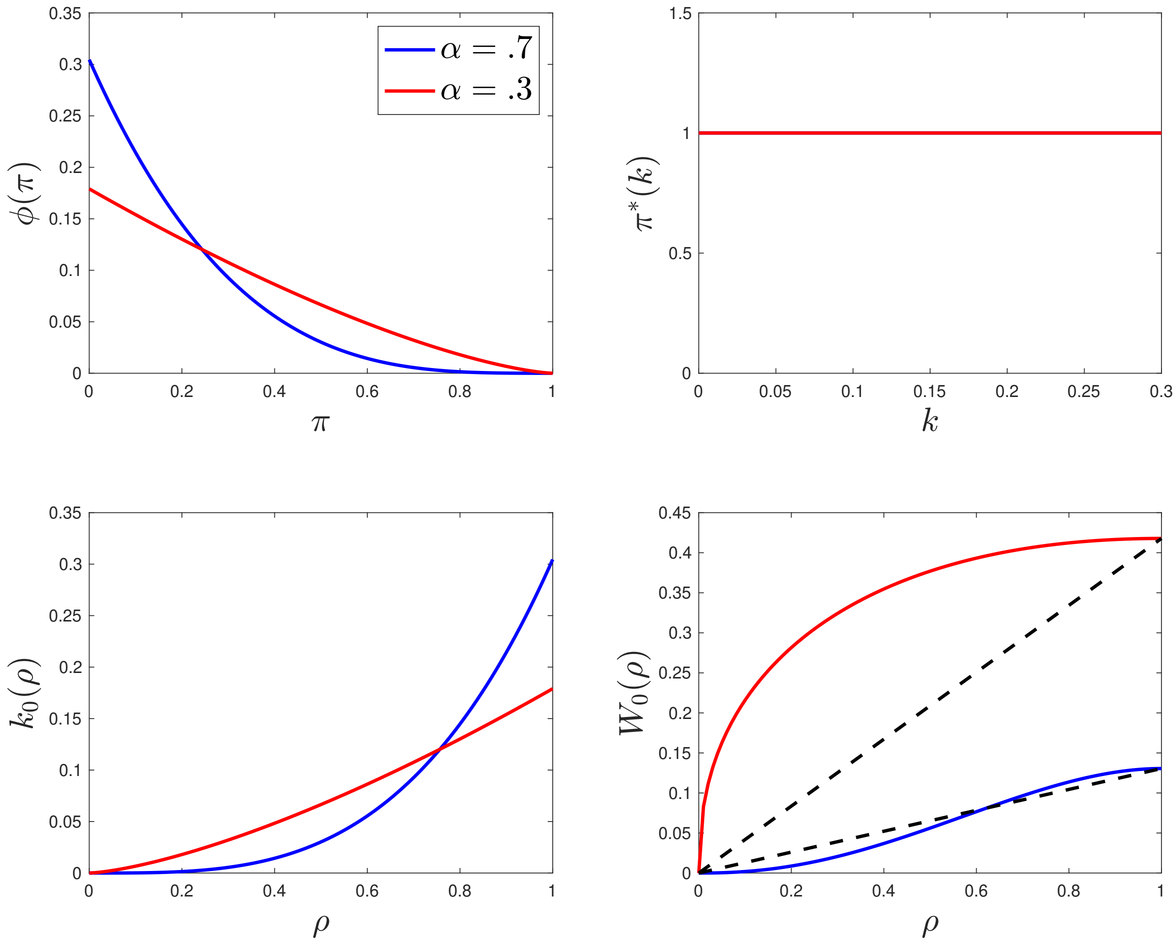

This example illustrates that whether W_0 is concave depends on how private agents respond to changes in expected policy. Intuitively, if private agents are cautious in that they need the policy to be sufficiently close to the Ramsey policy for them to take the good action then the complementarities between x and π are increasing in the level of reputation and W_0 may not be globally concave. Conversely, if agents are bold and a small change in the policy induces them to take an action close to the Ramsey policy when reputation is large then the complementarities between x and π are decreasing in the level of reputation, which helps generate a globally concave W_0.

C Problem (9) characterizes best outcome

In this section we show that the solution to problem (9) is the *best* Perfect Bayesian Equilibrium (PBE) outcome of the policy game described in the text. We will do this in two steps. First, we will show that the best PBE is the solution to a maximization problem similar to (9) that allows for randomization by the optimizing type. Second, we show that it is never optimal for the optimizing type to randomize and thus the problem in the first part reduces to (9) in the text.

C.1 Perfect Bayesian Equilibrium Outcome

We now define a PBE for the policy game. Let the horizon for the economy be $K \geqslant 1$ and let $t = 0, 1, ..., K-1$. (Note that there is an inverse relationship between periods and horizon so $t = K - k$. Here strategies and equilibrium objects will be indexed by t while in the text we go backward and index equilibrium objects by the residual horizon as it is more convenient.)

Let $h_r^t = (\pi_{r0}, x_0, \pi_0, ..., \pi_{rt-1}, x_{t-1}, \pi_{t-1}) \in H_r^t$ be the history faced by the rule designer, $h_x^t = \left(h_r^t, \pi_{rt}\right) \in H_x^t$ be the history faced by the private agents, and let $h_o^t = \left(h_x^t, x_t\right) \in H_o^t$ be the history faced by the optimizing type. A strategy for the rule designer is $\sigma_r = \{\sigma_{rt}\}$ where $\sigma_{rt} : H_r^t \to [\underline{\pi}, \bar{\pi}]$, an allocation rule for the private agents is $\sigma_x = \{\sigma_{xt}\}$ where $\sigma_{xt} : H_x^t \to X$, and a strategy for the optimizing type is $\sigma_o = \{\sigma_{ot}\}$ where $\sigma_{ot} : H_o^t \to \Delta\left([\underline{\pi}, \bar{\pi}]\right)$ and $\Delta\left([\underline{\pi}, \bar{\pi}]\right)$ is the space of probability measures over $[\underline{\pi}, \bar{\pi}]$. Beliefs are denoted by, $\rho = \{\rho_{t+1}\}$, where $\rho_t : H_r^t \to [0, 1]$.

Definition. Strategies $(\sigma_r, \sigma_x, \sigma_o)$ and beliefs ρ are a Perfect Bayesian Equilibrium (PBE) if:

1. For all histories h_r^t, the rule designer's strategy is optimal i.e. it maximizes the rule designer's expected payoff given ρ and (σ_o, σ_x);

2. For all histories h_o^t, the optimizing type's strategy is optimal i.e. it maximizes the optimizing type's expected payoff given ρ and (σ_r, σ_x);

3. The private agent's allocation rule satisfies

$$\sigma_x\left(h_r^t, \pi_{rt}\right) = \phi\left(\rho_t\left(h_r^t\right)\pi_{rt} + \left(1 - \rho_t\left(h_r^t\right)\right)\int \pi\sigma_o\left(\pi|h_o^t\right)d\pi\right)$$

   where $\sigma_{ot}\left(\pi_t|h_o^t\right)$ is the probability that σ_{ot} assigns to π_t after history h_o^t is reached;

4. Beliefs are updated using Bayes' rule wherever applicable, i.e. for all h_x^t along the equilibrium path,

$$\rho_{t+1}\left(h_r^{t+1}\right)=\frac{\rho_t\left(h_r^t\right)\mathbb{I}_{\{\pi_t=\pi_{rt}\}}}{\rho_t\left(h_r^t\right)\mathbb{I}_{\{\pi_t=\pi_{rt}\}}+\left(1-\rho_t\left(h_r^t\right)\right)\sigma_{ot}\left(\pi_t|h_o^t\right)}.$$

We next show that the problem below characterizes the best PBE:

$$W_{k+1}(\rho)=\max_{\pi_c,x,\sigma\in[0,1]}(\rho+(1-\rho)\sigma)\left[w(x,\pi_c)+\beta W_k\left(\frac{\rho}{\rho+(1-\rho)\sigma}\right)\right] \\ +(1-\rho)(1-\sigma)\left[w(x,\pi^*(x))+\beta W_k(0)\right] \tag{50}$$

subject to the implementability condition,

$$x=\phi\left((\rho+(1-\rho)\sigma)\pi_c+(1-\rho)(1-\sigma)\pi^*(x)\right), \tag{51}$$

and the incentive compatibility constraint for the optimizing type,

$$\sigma\left[w(x,\pi_c)+\beta_o V_k\left(\frac{\rho}{\rho+(1-\rho)\sigma}\right)-\left[w(x,\pi^*(x))+\beta_o V_k(0)\right]\right]=0. \tag{52}$$

The value for the optimizing type for a generic horizon $k+1$ is

$$V_{k+1}(\rho)=\sigma_k(\rho)\left[w\left(x_{k+1}(\rho),\pi_{c,k+1}(\rho)\right)+\beta_o V_k\left(\frac{\rho}{\rho+(1-\rho)\sigma}\right)\right] \\ +\left(1-\sigma_k(\rho)\right)\left[w\left(x_{k+1}(\rho),\pi^*\left(x_{k+1}(\rho)\right)\right)+\beta_o V_k(0)\right].$$

We prove the following result for the two period model but the logic can be extended inductively to arbitrary horizons.

Lemma 5. *The best PBE outcome solves the problem in (50).*

Proof. First, note that the solution to (50) can be supported as a PBE. Denote this solution with a hat. In period $t=0$, let the support of the optimizing type's strategy be $\{\pi_{rt},\pi^*(x_t)\}$. In particular, if $w(x_0,\pi_{r0})+\beta_o V_0(\rho)\geqslant w(x_0,\pi^*(x_0))+\beta_o V_0(0)$ set $\sigma_{o0}\left(\pi_{r0}|h_o^0\right)=1$. If instead $w(x_0,\pi_{r0})+\beta_o V_0(1)<w(x_0,\pi^*(x_0))+\beta_o V_0(0)$, set $\sigma_{o0}\left(\pi_{r0}|h_o^0\right)=0$. If none of these two conditions are satisfied, let $\sigma_{o0}\left(\pi_{r0}|h_o^0\right)=\tilde{\sigma}$ where $\tilde{\sigma}$ solves

$$w(x_0,\pi_{r0})+\beta_o V_0\left(\frac{\rho}{\rho+(1-\rho)\tilde{\sigma}}\right)=w(x_0,\pi^*(x_0))+\beta_o V_0(0).$$

In period 1, let $\sigma_{o1}\left(h_o^1\right)=\pi^*(x_1)$ for all ρ and x_1.

Let the evolution of beliefs be

$$\rho_1\left(h_r^1\right) = \begin{cases} \frac{\rho}{\rho+(1-\rho)\sigma_{o0}\left(\pi_{r0}|h_o^0\right)} & \text{if } \pi_0 = \pi_{r0} \\ 0 & \text{otherwise} \end{cases}$$

The strategy for the rule designer is

$$\sigma_{r0} = \hat{\pi}_{c1}$$

and

$$\sigma_{r1}\left(h_r^1\right) = \begin{cases} \pi_{c0}\left(\rho\right) & \text{if } \pi_0 = \pi_{r0} \text{ and } \sigma = 1 \\ \pi_{c0}\left(1\right) & \text{if } \pi_0 = \pi_{r0} \text{ and } \sigma = 0 \\ \pi_{c0}\left(0\right) & \text{if } \pi_0 \neq \pi_{r0} \end{cases}$$

where $\pi_{c0}\left(\rho\right)$ is defined in the text. Finally, let $\sigma_{xt}\left(h_x^t\right) = x_t$ where x_t solves

$$x_t = \phi\left(\left[\rho_t\left(h_r^t\right) + \left(1-\rho_t\left(h_r^t\right)\right)\sigma_{ot}\left(\pi_{rt}|\pi_{ot}h_o^t\right)\right]\pi_{rt} + \left(1-\rho_t\left(h_r^t\right)\right)\left(1-\sigma_{ot}\left(\pi_{rt}|\pi_{ot}h_o^t\right)\right)\pi^*\left(x_t\right)\right).$$

Clearly $\left(\sigma_r, \sigma_x, \sigma_o, \rho\right)$ constitutes a PBE that supports the outcome which solves (50).

Next, we show that no other PBE outcome can attain a higher value. Consider an arbitrary PBE. First, note that in the last period, given the prior ρ, there is a unique equilibrium outcome determined by the solution to problem (3). Thus, the continuation value for the rule designer and the optimizing type are $W_0\left(\rho\right)$ and $V_0\left(\rho\right)$ respectively.

Consider now the first period. We argue that the support of the optimal strategy for the optimizing type can contain only two possible elements: the rule and the best response to x_t. Suppose by way of contradiction that the optimizing type assigns probability to $\pi_t \neq \{\pi_{rt}, \pi^*\left(x_t\right)\}$. Since this is on path it must be that $\rho_{t+1} = 0$ so the value is $w\left(x_t, \pi_t\right) + \beta_o V_0\left(0\right) < w\left(x_t, \pi^*\left(x_t\right)\right) + \beta_o V_0\left(0\right)$. Thus it is better off choosing its static best response, a contradiction. The optimizing type's strategy can then be described by the probability of following the rule which we will denote by σ (with a slight abuse of notation). This immediately implies that the implementability constraint (51) must also be satisfied in any PBE outcome. There are three possible cases that can arise in equilibrium.

First, suppose that $\sigma \in (0,1)$. Since the optimizing type must be indifferent between following the rule and playing the static best response to x_0, it must be that

$$w\left(x_0, \pi_r\right) + \beta_o V_0\left(\frac{\rho}{\rho + (1-\rho)\sigma}\right) = w\left(x_0, \pi^*\left(x_0\right)\right) + \beta_o V_0\left(0\right)$$

or

$$\sigma\left[w\left(x_0, \pi_r\right) + \beta_o V_0\left(\frac{\rho}{\rho + (1-\rho)\sigma}\right) - \left[w\left(x_0, \pi^*\left(x_0\right)\right) + \beta_o V_0\left(0\right)\right]\right] = 0$$

Thus, the equilibrium outcome must satisfy all the constraints in (50). Hence $W_1(\rho)$ is (weakly) greater than the expected payoff of this PBE.

Second, suppose that $\sigma = 0$. This implies that

$$\sigma\left[w(x_0,\pi_r) + \beta_o V_0\left(\frac{\rho}{\rho+(1-\rho)\sigma}\right) - [w(x_0,\pi^*(x_0)) + \beta_o V_0(0)]\right] = 0$$

Thus, the equilibrium must satisfy all the constraints in (50). Hence $W_1(\rho)$ is (weakly) greater than the expected payoff of this PBE.

Finally, suppose $\sigma = 1$. Since the optimizing type always has the option to play the static best response to x_0, optimality requires

$$w(x_0,\pi_r) + \beta_o V_0(\rho) \geqslant w(x_0,\pi^*(x_0)) + \beta_o V_0(\rho_1(\pi^*(x_0)|\pi_{r0},x_0)) \tag{53}$$

where $\rho_1(\pi^*(x_0)|\pi_{r0},x_0) \geqslant 0$ is the posterior if the optimizing type deviates and $x_0 = \phi(\pi_{r0})$. Since $\sigma = 1$, a deviation only happens off-path and thus Bayes' rule does not pin down the posterior after a deviation. If $\rho_1(\pi^*(x_0)|\pi_{r0},x_0) = 0$ and (53) holds with equality then this PBE satisfies (52) and therefore it cannot attain a higher value then $W_1(\rho)$. Suppose now that either $\rho_1(\pi^*(x_0)|\pi_{r0},x_0) > 0$ or (53) holds with a strict inequality so

$$w(\phi(\pi_{r0}),\pi_{r0}) + \beta_o V_0(\rho) > w(\phi(\pi_{r0}),\pi^*(\phi(\pi_{r0}))) + \beta_o V_0(0) \tag{54}$$

since V_0 is strictly decreasing. Next we show that in this case there is an outcome feasible in (50) that attains a higher value. In particular, the rule designer can then reduce π_{r0} by $\varepsilon > 0$ until the constraint holds with equality,

$$w(\phi(\pi_{r0}-\varepsilon),\pi_{r0}-\varepsilon) + \beta_o V_0(\rho) = [w(\phi(\pi_{r0}-\varepsilon),\pi^*(x_0-\varepsilon)) + \beta_o V_0(0)].$$

Since the solution must be bounded away from the Ramsey outcome because of Assumption 3, this perturbation increases welfare and it is feasible in (50).

Thus the solution to the problem in (50) is the PBE outcome that attains the highest value. Q.E.D.

C.2 Optimizing type does not randomize

We now show that under our assumptions it is without loss of generality to consider the case in which the optimizing type either follows the rule with probability one or chooses its best response and deviates from the rule with probability one.

To see this, note that we can write problem (50) where we allow the optimizing type to randomize as

$$W_1(\rho) = \max\left\{\max_{\sigma\in[0,1]} W_{pool}(\sigma,\rho), W_{sep}(\rho)\right\}$$

where $W_{sep}(\rho)$ is the value of separation – defined in the text – and $W_{pool}(\sigma,\rho)$ is the value the rule designer can attain by inducing the optimizing type to follow the rule with probability σ starting with a prior ρ is

$$\begin{aligned} W_{pool}(\sigma,\rho) = {} & [\rho + (1-\rho)\sigma]\left[w\left(x_{ico}(\sigma,\rho), \pi_{ico}(\sigma,\rho)\right) + \beta W_0\left(\rho'(\rho,\sigma)\right)\right] \\ & + (1-\rho)(1-\sigma)\left[w\left(x_{ico}(\sigma,\rho), \pi^*\left(x_{ico}(\sigma,\rho)\right)\right) + \beta W_0(0)\right] \end{aligned}$$

where the evolution of the prior is given by

$$\rho'(\sigma,\rho) = \frac{\rho}{\rho + (1-\rho)\sigma}$$

and $\left(x_{ico}(\sigma,\rho), \pi_{ico}(\sigma,\rho)\right)$ solves

$$x_{ico}(\sigma,\rho) = \phi\left([\rho + (1-\rho)\sigma]\,\pi_{ico}(\sigma,\rho) + (1-\rho)(1-\sigma)\,\pi^*\left(x_{ico}(\sigma,\rho)\right)\right)$$

and

$$w\left(x_{ico}(\sigma,\rho), \pi_{ico}(\sigma,\rho)\right) + \beta V_0\left(\rho'(\rho,\sigma)\right) = w\left(x_{ico}(\sigma,\rho), \pi^*\left(x_{ico}(\sigma,\rho)\right)\right) + \beta V_0(0).$$

That is, $\pi_{ico}(\sigma,\rho)$ is the most stringent policy that is incentive compatible for the optimizing type given ρ and σ.

We next show that $W_{pool}(1,\rho) \geqslant W_{pool}(\sigma,\rho)$ thus it is optimal to choose $\sigma \in \{0,1\}$. We will use the following intermediate result:

Lemma 6. *Under Assumptions 1 and 2, $V_0(\rho)$ is concave.*

Proof. We prove this for the case with $w_x < 0$. The proof for the other case is identical. Recall that

$$V_0(\rho) = w\left(x(\rho), \pi^*\left(x(\rho)\right)\right)$$

Then

$$\begin{aligned} V_0'(\rho) &= w_x x'(\rho) + w_\pi \pi_x^*(x)\, x'(\rho) \\ &= w_x x'(\rho) \end{aligned}$$

where we used that $w_\pi(x, \pi^*(x)) = 0$. Then

$$\begin{aligned} V''(\rho) &= (w_{xx} + w_{x\pi}\pi^*_x)\, x'(\rho)^2 + w_x x''(\rho) \\ &= \left(\frac{w_{xx} w_{\pi\pi}(x, \pi^*) - w_{x\pi}(x, \pi^*)^2}{w_{\pi\pi}(x, \pi_o)} \right) x'(\rho)^2 + w_x x''(\rho) \\ &< 0 \end{aligned}$$

where the second line follows from using (21) to substitute for π^*_x and the last inequality follows from Assumption 2 and $x''(\rho) \geqslant 0$ where the latter property was established in the proof of Proposition 1. Q.E.D.

Lemma 7. *Under Assumptions 1 and 2, for all ρ and σ, $W_{pool}(1, \rho) \geqslant W_{pool}(\sigma, \rho)$.*

Proof. Consider $\sigma < 1$. Consider a deviation in which the optimizing type chooses the following policy

$$\pi_{dev} = [\rho + (1-\rho)\sigma]\, \pi_{ico}(\sigma, \rho) + (1-\rho)(1-\sigma)\, \pi^*(x_{ico}(\sigma, \rho))$$

with probability one and $\pi_c = \pi_{dev}$. Note that this policy is just the expected value of the policies. Therefore, under this deviation, $x_{dev} = x_{ico}(\sigma, \rho)$. Since w is concave in π and W_0 is concave in ρ, this policy improves welfare:

$$\begin{aligned} W_{pool}(\sigma, \rho) &= [\rho + (1-\rho)\sigma]\, w(x_{ico}(\sigma, \rho), \pi_{ico}(\sigma, \rho)) + (1-\rho)(1-\sigma)\, w(x_{ico}(\sigma, \rho), \pi^*(x_{ico}(\sigma, \rho))) \\ &\quad + [\rho + (1-\rho)\sigma]\, \beta W_0\left(\rho'(\rho, \sigma)\right) + (1-\rho)(1-\sigma)\, \beta W_0(0) \\ &\leqslant w(x_{ico}(\sigma, \rho), [\rho + (1-\rho)\sigma]\, \pi_{ico}(\sigma, \rho) + (1-\rho)(1-\sigma)\, \pi^*(x_{ico}(\sigma, \rho))) + \beta W_0(\rho) \\ &= w(x_{ico}(\sigma, \rho), \pi_{dev}) + \beta W_0(\rho) \end{aligned}$$

We are left to show that this deviation is feasible for the rule designer in that it satisfies the incentive compatibility constraint for the optimizing type:

$$w(x_{ico}(\sigma, \rho), \pi_{dev}) + \beta V_0(\rho) \geqslant w(x_{ico}(\sigma, \rho), \pi^*(x_{ico}(\sigma, \rho))) + \beta V_0(0)$$

Note that at the original allocation it must be that

$$w(x_{ico}(\sigma, \rho), \pi_{ico}(\sigma, \rho)) + \beta V_0\left(\rho'(\rho, \sigma)\right) \geqslant w(x_{ico}(\sigma, \rho), \pi^*(x_{ico}(\sigma, \rho))) + \beta V_0(0) \quad (55)$$

and trivially

$$w(x_{ico}(\sigma, \rho), \pi^*(x_{ico}(\sigma, \rho))) + \beta V_0(0) \geqslant w(x_{ico}(\sigma, \rho), \pi^*(x_{ico}(\sigma, \rho))) + \beta V_0(0). \quad (56)$$

Multiplying the left and right side of (55) by $[\rho + (1-\rho)\sigma]$, the left and right side of (56) by $(1-\rho)(1-\sigma)$, and summing up the two resulting equations yields

$$\begin{aligned} & w\left(x_{ico}(\sigma,\rho), \pi^*\left(x_{ico}(\sigma,\rho)\right)\right) + \beta V_0(0) \\ & \leqslant [\rho + (1-\rho)\sigma]\left[w\left(x_{ico}(\sigma,\rho), \pi_{ico}(\sigma,\rho)\right) + \beta V_0\left(\rho'(\rho,\sigma)\right)\right] \\ & + (1-\rho)(1-\sigma)\left[w\left(x_{ico}(\sigma,\rho), \pi^*\left(x_{ico}(\sigma,\rho)\right)\right) + \beta V_0(0)\right] \\ & \leqslant w\left(x_{ico}(\sigma,\rho), \pi_{dev}\right) + \beta V_0(\rho) \end{aligned} \tag{57}$$

where the second inequality follows from concavity of w in π and V_0 in ρ. Thus the proposed deviation is incentive compatible and it increases welfare. Moreover, since

$$W_{pool}(1,\rho) = \max_{\pi_c} w\left(\phi(\pi_c), \pi_c\right) + \beta W_0(\rho)$$

subject to

$$w\left(\phi(\pi_c), \pi_c\right) + \beta V_0(\rho) \geqslant w\left(\phi(\pi_c), \pi^*\left(\phi(\pi_c)\right)\right) + \beta V_0(0)$$

and since π_{dev} is feasible for this problem we have

$$W_{pool}(1,\rho) \geqslant w\left(x_{ico}(\sigma,\rho), \pi_{dev}\right) + \beta W_0(\rho) \geqslant W_{pool}(\sigma,\rho)$$

as wanted. Q.E.D.

This lemma immediately implies that

$$W_1(\rho) = \max\left\{W_{pool}(1,\rho), W_{sep}(\rho)\right\} = \max\left\{W_{pool}(\rho), W_{sep}(\rho)\right\}$$

as defined in the text.

D Proof of Proposition 2 for the Barro-Gordon example

Consider the Barro-Gordon example with

$$\phi\left(\rho\pi_c + (1-\rho)\pi_o\right) = \rho\pi_c + (1-\rho)\pi_o$$

and

$$w(x,\pi) = -\frac{1}{2}\left[(\psi + x - \pi)^2 + \pi^2\right].$$

Therefore,

$$w_\pi = -\left[-(\psi + x) + 2\pi\right],$$

$$w_x = -(\psi + x - \pi).$$

Moreover, the static best response is

$$\pi^*(x) = \frac{(\psi + x)}{2}.$$

Consider first the static problem:

$$W_0(\rho) = \max_{\pi_c, \pi_o, x} -\rho\frac{1}{2}\left[(\psi + x - \pi_c)^2 + \pi_c^2\right] - (1-\rho)\frac{1}{2}\left[(\psi + x - \pi_o)^2 + \pi_o^2\right]$$

subject to

$$x = \rho\pi_c + (1-\rho)\pi_o,$$

$$\pi_o = \frac{(\psi + x)}{2}.$$

Combining the two constraints we can express π_o and x in terms of π_c as

$$\pi_o = \frac{\psi + \rho\pi_c}{(1+\rho)} = \frac{\rho}{1+\rho}\pi_c + \frac{\psi}{(1+\rho)}$$

and

$$x = \frac{2\rho}{1+\rho}\pi_c + \frac{(1-\rho)}{(1+\rho)}\psi.$$

Therefore, substituting into the objective function we obtain

$$\begin{aligned} W_0(\rho) = \max -\frac{1}{2}\Bigg(\rho\left[\left(\psi + \frac{2\rho}{1+\rho}\pi_c + \frac{(1-\rho)}{(1+\rho)}\psi - \pi_c\right)^2 + \pi_c^2\right] \\ + (1-\rho)\left[\left(\psi + \frac{2\rho}{1+\rho}\pi_c + \frac{(1-\rho)}{(1+\rho)}\psi - \frac{\rho}{1+\rho}\pi_c - \frac{\psi}{(1+\rho)}\right)^2 + \left(\frac{\rho}{1+\rho}\pi_c + \frac{\psi}{(1+\rho)}\right)^2\right]\Bigg) \\ = \max_{\pi_c} -\frac{1}{2}\left(\rho\left[\left(\frac{2}{(1+\rho)}\psi - \left(\frac{1-\rho}{1+\rho}\right)\pi_c\right)^2 + \pi_c^2\right] + (1-\rho)\left[2\left(\frac{\psi}{(1+\rho)} + \frac{\rho}{1+\rho}\pi_c\right)^2\right]\right). \end{aligned}$$

The first order condition for the above problem is

$$\begin{aligned} 0 = \left[-\left(\left(\frac{1-\rho}{1+\rho}\right)\frac{4\rho}{(1+\rho)}\psi - \rho\left(\frac{1-\rho}{1+\rho}\right)2\left(\frac{1-\rho}{1+\rho}\right)\pi_c\right) + 2\rho\pi_c\right] \\ + \left[4\frac{\rho(1-\rho)}{1+\rho}\left(\frac{\psi}{(1+\rho)} + \frac{\rho}{1+\rho}\pi_c\right)\right] \end{aligned}$$

which implies that

$$\pi_c = 0$$

Therefore

$$\pi_o(\rho) = \frac{\psi}{(1+\rho)},$$

$$x(\rho) = \frac{(1-\rho)}{(1+\rho)}\psi$$

and

$$W_0(\rho) = -\frac{\psi^2}{(1+\rho)}.$$

Also it is worth noting that

$$W_0''(\rho) = -2\frac{\psi^2}{(1+\rho)} < 0$$

and so W_0 is concave.

Consider now the two-period problem. Let the value for the optimizing type in the terminal period be

$$\begin{aligned} V(\rho) &= -\frac{1}{2}\left[(\psi + x(\rho) - \pi^*)^2 + (\pi^*(x(\rho)))^2\right] \\ &= -\left(\frac{\psi}{(1+\rho)}\right)^2. \end{aligned}$$

Lets consider the pooling case first. The optimal rule solves

$$W^{pool}(\rho) = \max_{\pi,x} -\frac{1}{2}\left[(\psi + x - \pi)^2 + \pi^2\right] + \beta W_0(\rho)$$

subject to

$$x = \pi$$

$$-\frac{1}{2}\left[(\psi + x - \pi)^2 + \pi^2\right] + \beta_o V(\rho) \geqslant -\frac{1}{2}\left[(\psi + x - \pi^*(x))^2 + \pi^*(x)^2\right] + \beta_o V(0)$$

or

$$W^{pool}(\rho) = \max_{\pi} -\frac{1}{2}\left[\psi^2 + \pi^2\right] + \beta W_0(\rho)$$

subject to

$$-\frac{1}{2}\left[\psi^2 + \pi^2\right] - \beta_o\left(\frac{\psi}{(1+\rho)}\right)^2 = -\frac{(\psi+\pi)^2}{4} - \beta_o \psi^2.$$

Thus the solution is pinned down by the last constraint. Solving for π we have

$$\pi_{ico}(\rho) = \psi\left(1 - \sqrt{4\beta_o\left[1 - \left(\frac{1}{(1+\rho)}\right)^2\right]}\right)$$

(Note that $\pi_{ico}(0) = \psi = \pi_o(0)$ and $\pi_{ico}(1) = \psi\left(1 - \sqrt{3\beta_o}\right)$ so it must be that $\beta_o < 1/3$.)

So the payoff from pooling is

$$W^{pool}(\rho) = -\psi^2 \left[\frac{1}{2} \left[1 + \left(1 - \sqrt{4\beta_o \left[1 - \left(\frac{1}{(1+\rho)} \right)^2 \right]} \right)^2 \right] + \beta \frac{1}{(1+\rho)} \right].$$

The value of separation is

$$\begin{aligned} W^{sep}(\rho) &= W_0(\rho) + \beta \left[\rho W_0(1) + (1-\rho) W_0(0) \right] \\ &= -\psi^2 \left[\frac{1}{(1+\rho)} + \beta \left[1 - \frac{\rho}{2} \right] \right]. \end{aligned}$$

Let's consider

$$\begin{aligned} \Delta(\rho) \equiv & W^{pool}(\rho) - W^{sep}(\rho) \\ = & \psi^2 \left(- \left[\frac{1}{2} \left[1 + \left(1 - \sqrt{4\beta_o \left[1 - \left(\frac{1}{(1+\rho)} \right)^2 \right]} \right)^2 \right] + \beta \frac{1}{(1+\rho)} \right] \right. \\ & \left. + \left[\frac{1}{(1+\rho)} + \beta \left[1 - \frac{\rho}{2} \right] \right] \right) \\ = & \psi^2 \left(\left[\frac{1}{(1+\rho)} - \frac{1}{2} \left[1 + z(\rho) \right] \right] + \beta \left[1 - \frac{\rho}{2} - \frac{1}{1+\rho} \right] \right) \end{aligned}$$

where

$$z(\rho) \equiv \left(1 - \sqrt{4\beta_o \left[1 - \left(\frac{1}{(1+\rho)} \right)^2 \right]} \right)^2 = \left(1 - \frac{2\xi}{(1+\rho)} \sqrt{\left[(1+\rho)^2 - 1 \right]} \right)^2 < 1$$

and

$$\xi \equiv \sqrt{\beta_o}.$$

At $\rho = 0$, $\Delta(0) = 0$, while at $\rho = 1$,

$$\Delta(1) = \psi^2 \left(\left[-\frac{1}{2} \left(1 - \sqrt{4\beta_o \left[1 - \left(\frac{1}{2} \right)^2 \right]} \right)^2 \right] \right) < 0.$$

Next, let's look at the slope

$$\Delta'(\rho) = \psi^2 \left(\left[-\frac{1}{(1+\rho)^2} - \frac{1}{2} z'(\rho) \right] + \beta \left[-\frac{1}{2} + \frac{1}{(1+\rho)^2} \right] \right)$$

$$z'(\rho) = -\left(8\beta_o \frac{1}{(1+\rho)^3}\right)\left(\left(4\beta_o\left[1-\left(\frac{1}{(1+\rho)}\right)^2\right]\right)^{-\frac{1}{2}} - 1\right) \tag{58}$$

Notice that as $\rho \to 0$, $z'(\rho)$ goes to $-\infty$. Therefore the slope of $\Delta(\rho)$ at $\rho = 0$ is

$$\psi^2\left(\left[-1-\frac{1}{2}z'(\rho)\right]+\frac{\beta}{2}\right) \to \infty$$

so that near 0 there exists a region of pooling.

In general to get pooling we need

$$\beta > \frac{\left[-\frac{1}{(1+\rho)}+\frac{1}{2}\left[1+z(\rho)\right]\right]}{\left[1-\frac{\rho}{2}-\frac{1}{1+\rho}\right]}.$$

Let

$$F(\rho) \equiv \frac{\left[-\frac{1}{(1+\rho)}+\frac{1}{2}\left[1+z(\rho)\right]\right]}{\left[1-\frac{\rho}{2}-\frac{1}{1+\rho}\right]} \tag{59}$$

We know that $\lim_{\rho\to 0} F(\rho) = -\infty$ and $\lim_{\rho\to 1} F(\rho) = \infty$. In fact, we have

$$F(0) = \frac{0}{0}$$

Thus, by the L'Hôpital's rule we have

$$\lim_{\rho\to 0} F(0) = \lim_{\rho\to 0} \frac{\left[1+\frac{1}{2}z'(0)\right]}{\frac{1}{2}} = -\infty$$

and

$$F(1) = \frac{\left[-\frac{1}{2}+\frac{1}{2}\left[1+z(1)\right]\right]}{\left[\frac{1}{2}-\frac{1}{2}\right]} = \infty$$

We next show that $F(\rho)$ is monotone increasing in ρ so that there exists a cutoff ρ^* such that it is optimal to pool for $\rho < \rho^*$ and it is optimal to separate for $\rho > \rho^*$. To this end, note that we can rearrange (59) as

$$F(\rho) = \frac{-2+(1+\rho)\left[1+z(\rho)\right]}{\rho(1-\rho)}.$$

So

$$F'(\rho) = \frac{\rho(1-\rho)\left[(1+\rho)z'(\rho)+1+z(\rho)\right]-\left[-2+(1+\rho)\left[1+z(\rho)\right]\right]\left[1-2\rho\right]}{\rho^2(1-\rho)^2}. \tag{60}$$

The denominator is positive and so we just need to sign the numerator in order to sign $F'(\rho)$.

Lets do some preliminary calculations. We have that

$$\sqrt{z(\rho)} = 1 - 2\xi \frac{\sqrt{\left[(1+\rho)^2 - 1\right]}}{(1+\rho)}$$

and

$$\frac{2\xi}{(1+\rho)} = \frac{1-\sqrt{z(\rho)}}{\sqrt{\left[(1+\rho)^2-1\right]}}. \tag{61}$$

Therefore, after some algebraic manipulations, we can rewrite (58) as

$$\begin{aligned} z'(\rho) &= -\left(8\beta_o \frac{1}{(1+\rho)^3}\right)\left(\left(4\beta_o\left[1-\left(\frac{1}{(1+\rho)}\right)^2\right]\right)^{-\frac{1}{2}} - 1\right) \\ &= -\left(\frac{2}{(1+\rho)}\left(\frac{2\xi}{(1+\rho)}\right)^2\right)\left(\frac{1}{2\xi\frac{\sqrt{[(1+\rho)^2-1]}}{(1+\rho)}} - 1\right) \\ &= -\frac{2}{(1+\rho)}\frac{\left(\sqrt{z(\rho)} - z(\rho)\right)}{\rho(2+\rho)} \end{aligned} \tag{62}$$

where in the third line we used (61).

Next, let's consider the numerator of $F'(\rho)$ in (60):

$$\begin{aligned} h(\rho) &\equiv \rho(1-\rho)\left[(1+\rho)z'(\rho) + 1 + z(\rho)\right] - \left[-2 + (1+\rho)\left[1+z(\rho)\right]\right][1-2\rho] \\ &= (1-\rho)\left[\frac{-2\sqrt{z(\rho)} + (2+\rho(2+\rho))z(\rho) + \rho(2+\rho)}{(2+\rho)}\right] \\ &\quad -\left\{-2[1-2\rho] + (1+\rho)[1-2\rho] + (1+\rho)[1-2\rho]z(\rho)\right\} \\ &= \left[\frac{-2(1-\rho)\sqrt{z(\rho)} + \left[\rho+4\rho^2+\rho^3\right]z(\rho) + (2+\rho)(1-\rho)^2}{(2+\rho)}\right] \\ &> \left[\frac{-2(1-\rho)\sqrt{z(\rho)} + \left[\rho+4\rho^2\right]z(\rho) + (2+\rho)(1-\rho)^2}{(2+\rho)}\right] \\ &= \left[\frac{(1-\rho)\left[(2+\rho)(1-\rho) - 2\sqrt{z(\rho)}\right] + \left[\rho+4\rho^2\right]z(\rho)}{(2+\rho)}\right] \end{aligned}$$

where to obtain the second equality we used (62), the third equality follows from algebra, the inequality follows from $\rho^3 z(\rho) > 0$, and the last equality also follows from algebra. Notice that if $\left[(2+\rho)(1-\rho) - 2\sqrt{z(\rho)}\right] \geqslant 0$ then $h(\rho) > 0$ and the result is proved. Suppose not, i.e.

$$(1-\rho) < \frac{2\sqrt{z(\rho)}}{(2+\rho)} \tag{63}$$

We can then write

$$\begin{aligned}
h(\rho) &> \left[\frac{(1-\rho)\left[(2+\rho)(1-\rho) - 2\sqrt{z(\rho)}\right] + \left[\rho + 4\rho^2\right] z(\rho)}{(2+\rho)}\right] \\
&> \left[\frac{\frac{2\sqrt{z(\rho)}}{(2+\rho)}\left[(2+\rho)(1-\rho) - 2\sqrt{z(\rho)}\right] + \left[\rho + 4\rho^2\right] z(\rho)}{(2+\rho)}\right] \\
&= \left[\frac{2(1-\rho)\sqrt{z(\rho)} - z(\rho)\left[\frac{4}{(2+\rho)} - \rho - 4\rho^2\right]}{(2+\rho)}\right] \\
&> \frac{z(\rho)}{(2+\rho)}\left[2(1-\rho) - \left[\frac{4}{(2+\rho)} - \rho - 4\rho^2\right]\right]
\end{aligned}$$

where the second line follows from (63), the third is algebra, and the fourth line follows from $\sqrt{z(\rho)} > z(\rho)$ since $z(\rho) < 1$. We next show that $2(1-\rho) > \left[\frac{4}{(2+\rho)} - \rho - 4\rho^2\right]$. To see this, suppose it is not true. Then

$$2(1-\rho) - \left[\frac{4}{(2+\rho)} - \rho - 4\rho^2\right] < 0$$

or

$$4\rho^2(2+\rho) - \rho^2 < 0$$

which is a contradiction. Therefore $h(\rho) > 0$ and so $F'(\rho) > 0$. Q.E.D.

E Optimal rules when the rule designer can commit

In our baseline model, we assumed that the rule designer chooses the optimal rule in each period without commitment. We now study the problem for a rule designer who can choose rules for all subsequent periods in period zero and commit to them. First, we show that in the twice repeated economy, the solution with and without commitment on the part of the rule designer coincides. With more than two periods, whether these two values coincide depends on the level of reputation. In particular, for reputation values that are either sufficiently high or low, the commitment and no-commitment outcomes coincide and so our main results extend to the setting where the rule designer has commitment. However, for an intermediate range of priors, the solution to this problem differs from the case in which rules are chosen sequentially: the rule designer itself suffers from a time inconsistency problem. This is because future rules can be used to incentivize the policy maker in the current period, thereby relaxing the incentive compatibility constraint. We illustrate this point in the simplest possible way by considering a thrice repeated economy.

Setup and preliminaries For a generic horizon k, we can write the rule designer's problem in a recursive fashion by ensuring that the continuation value delivers a given promised value to the optimizing type given the prior ρ.

To set up the problem, let $\mathcal{V}_{k+1}(\rho)$ be the set of feasible promised values for the optimizing type. Note that whenever $\rho = 0$, $\mathcal{V}_k(\rho) = \{V_k(0)\}$ for all k because π_r does not affect the equilibrium outcome which is given by $x = \phi(\pi^*(x))$ and $\pi_o = \pi^*(x)$. This observation implies that the worst punishment that the rule designer can impose on a deviating optimizing type is $V_k(0)$ because after a deviation private agents learn that they are facing the optimizing type. Thus, for $k \geqslant 1$ we can write

$$\begin{aligned}
\mathcal{V}_k(\rho) = \{V : \exists \left(x, \pi_r, \sigma, V', \rho'\right) & \text{ such that} \\
& V = \sigma\left[w(x,\pi_r) + \beta_o V'\right] + (1-\sigma)\left[w(x,\pi^*(x)) + \beta_o V_k(0)\right] \\
& x = \phi\left((\rho + \sigma(1-\rho))\pi_r + (1-\sigma)(1-\rho)\pi^*(x)\right) \\
& 0 = \sigma\left\{w(x,\pi_r) + \beta_o V' - \left[w(x,\pi^*(x)) + \beta_o V_k(0)\right]\right\} \text{ if } \sigma < 1 \\
& w(x,\pi_r) + \beta_o V' \geqslant w(x,\pi^*(x)) + \beta_o V_k(0) \text{ if } \sigma = 1 \\
& \rho' = \frac{\rho}{\rho + (1-\rho)\sigma} \\
& V' \in \mathcal{V}_{k-1}\left(\rho'\right)\}
\end{aligned}$$

and

$$\mathcal{V}_0(\rho) = \{V : \exists (x, \pi_c) \text{ such that } V = w(x, \pi^*(x)) \text{ and } x = \phi(\rho\pi_c + (1-\rho)\pi^*(x))\}.$$

In setting up these values we used the fact that after any deviation the policy maker's continuation value is $V_k(0)$ defined in the text.

The problem of the rule designer with horizon $k \geqslant 1$ given promised value V and prior ρ is

$$\begin{aligned}
\bar{W}_k(V,\rho) = \max_{x,\pi_r,\sigma,V',\rho'} & (\rho + \sigma(1-\rho))\left[w(x,\pi_r) + \beta \bar{W}_{k-1}\left(V',\rho'\right)\right] \\
& + (1-\sigma)(1-\rho)\left[w(x,\pi^*(x)) + \beta W_{k-1}(0)\right]
\end{aligned} \tag{64}$$

subject to

$$\begin{aligned} V =& \sigma\left[w(x,\pi_r)+\beta_o V'\right]+(1-\sigma)\left[w(x,\pi^*(x))+\beta_o V_{k-1}(0)\right] \\ x =& \phi\left((\rho+\sigma(1-\rho))\pi_r+(1-\sigma)(1-\rho)\pi^*(x)\right) \\ 0 =& \sigma\left[w(x,\pi_r)+\beta_o V'-(w(x,\pi^*(x))+\beta_o V_{k-1}(0))\right] \\ \rho' =& \frac{\rho}{\rho+(1-\rho)\sigma} \\ V' \in& \mathcal{V}_{k-1}(\rho') \end{aligned}$$

where we used the fact that after a deviation by the optimizing type the value for the policy maker is $W_{k-1}(0)=\bar{W}_{k-1}(V_{k-1}(0),0)$. (Note that at an optimum, even if $\sigma=1$, the incentive constraint will be binding so it is without loss of generality to write the incentive constraint as an equality.)

The problem in the first period is the same without the promise keeping constraint or

$$\bar{W}_K(\rho)=\max_{V\in\mathcal{V}_K(\rho)}\bar{W}_K(V,\rho).$$

E.1 Two period problem

Proposition 8. *In the twice repeated economy, the solution with and without commitment on the part of the rule designer coincides.*

Proof. First, note that there is no disagreement between the rule designer and the optimizing type policy maker in the choice of the policy rule in a static setting. That is, both the rule designer's and policy maker's continuation values are maximized by setting $\pi_c=\underline{\pi}$ in the terminal period, or:

$$\max_{V\in\mathcal{V}_0(\rho)} V=V_0(\rho) \tag{65}$$

and

$$\max_{V\in\mathcal{V}_0(\rho)}\bar{W}_0(V,\rho)=\bar{W}(V_0(\rho),\rho)=W_0(\rho). \tag{66}$$

Condition (66) follows because $w(x,\pi^*(x))$ is decreasing in x thus choosing $\pi_c=\underline{\pi}$ maximizes both the value of the rule designer and the optimizing type.[19] Condition (65) follows from Proposition 1 that characterizes the relaxed problem without the promise keeping constraint. Thus, promising $V_0(\rho)$ maximizes the continuation value and at the same time relaxes the IC in the first period. Thus, the problem in (64) for $k=1$ dropping

[19]Note that this is no longer true if there are more than two periods.

the promise keeping constraint (since there are no promises in the first period) reduces to

$$\bar{W}_1(\rho) = \max_{x,\pi_r,\sigma,\rho'} (\rho + \sigma(1-\rho))\left[w(x,\pi_r) + \beta W_0(\rho')\right] + (1-\sigma)(1-\rho)\left[w(x,\pi^*(x)) + \beta W_0(0)\right]$$

subject to

$$\begin{aligned} x =& \phi\left((\rho + \sigma(1-\rho))\pi_r + (1-\sigma)(1-\rho)\pi^*(x)\right) \\ 0 =& \sigma\left[w(x,\pi_r) + \beta_o V_0(\rho') - (w(x,\pi^*(x)) + \beta_o V_0(0))\right] \\ \rho' =& \frac{\rho}{\rho + (1-\rho)\sigma} \end{aligned}$$

which is the same problem as the one where the rule designer lacks commitment. Q.E.D.

E.2 Three period problem

Proposition 9. *When* $K = 3$*:*

1. *For* $\rho \in [0, \hat{\rho}]$*, the optimal rules with and without commitment coincide and call for pooling in period 0 and 1.*

2. *For* $\rho \in [\rho_2^*, 1]$*, the optimal rules with and without commitment coincide and call for separation in period 0 and 1.*

3. *There exists* $\rho_3^* \in (\hat{\rho}, \rho_2^*)$ *such that for* $\rho \in (\rho_3^*, \rho_2^*)$ *the optimal rules with and without commitment do not coincide. In particular, without commitment, it is optimal to separate in period zero. With commitment, it is optimal to pool in period zero and commit to separation in period one. This is achieved by committing to the most stringent rule,* $\underline{\pi}$*, in period one.*

Proof. We start with an intermediate result. For $\rho \in [0, \hat{\rho}] \cup [\rho_2^*, 1]$ we have that

$$\max_{V \in \mathcal{V}_1(\rho)} V = V_1(\rho) \tag{67}$$

and

$$\max_{V \in \mathcal{V}_1(\rho)} \bar{W}_1(V, \rho) = \bar{W}_1(V_1(\rho), \rho) = W_1(\rho). \tag{68}$$

This follows from the argument provided in part 2 and 3 of Proposition 3. Instead, for $\rho \in (\hat{\rho}, \rho_2^*)$, the rule designer and the policy maker do not agree on the continuation outcome. In particular, for ρ in this range, the value for the optimizing type policy maker is maximized by having separation in the second period ($k = 1$) because $x_0(\rho) = \rho\underline{\pi} +$

$(1-\rho)\,\pi^*(x_0(\rho)) < \pi_{ico,1}(\rho)$ and

$$\max_{V\in\mathcal{V}_1(\rho)} V = V_1^{sep}(\rho) = w\left(x_0(\rho), \pi^*(x_0(\rho))\right) + \beta_o V_0(0)$$

However, there exists $\rho_3 \in \left[\hat{\rho}, \rho_2^*\right]$ such that for $\rho \in \left(\rho_3^*, \rho_2^*\right)$ it is optimal to pool for the rule designer. Note that by Proposition 2 we know that for $\rho \in \left(\hat{\rho}, \rho_1^*\right)$ it is optimal to pool and for $\rho \in \left[\rho_2^*, 1\right]$ it is optimal to separate. If there is a single cutoff as in the Barro-Gordon model, then $\rho_1^* = \rho_2^*$ and $\rho_3^* = \hat{\rho}$. If $\rho_1^* \neq \rho_2^*$ there can be multiple cutoffs in the region $\left(\rho_1^*, \rho_2^*\right)$ but by the definition of ρ_2^* we know that it is optimal to pool in an interval to the left of ρ_2^*, $\left(\rho_3^*, \rho_2^*\right) = \left(\rho_2^* - \varepsilon, \rho_2^*\right)$ for some $\varepsilon > 0$ sufficiently small. Thus for $\rho \in \left(\rho_3^*, \rho_2^*\right)$,

$$\max_{V\in\mathcal{V}_1(\rho)} W_1(V,\rho) = W_1(\rho) = w\left(x_{ico,1}(\rho), \pi_{ico,1}(\rho)\right) + \beta W_0(\rho) \neq \bar{W}_1\left(V_1^{sep}(\rho), \rho\right).$$

See Proposition 3 and Figure 5 for an illustration.

Part 1 and 2. Take $\rho \in [0,\hat{\rho}] \cup \left[\rho_2^*, 1\right]$. By the same argument provided in part 2 and 3 of Proposition 3, randomization is not optimal. Thus, since with $\sigma \in \{0,1\}$ it follows that $\rho' \in \{0, \rho, 1\}$ where (67) and (68) hold, the problem in (64) for $k = 2$ dropping the promise keeping constraint (since there are no promises in the first period) reduces to

$$\max_{x,\pi_r,\rho',\sigma\in\{0,1\}} (\rho + \sigma(1-\rho))\left[w(x,\pi_r) + \beta W_1\left(\rho'\right)\right] + (1-\sigma)(1-\rho)\left[w(x,\pi^*(x)) + \beta W_1(0)\right]$$

subject to

$$\begin{aligned} x =& \phi\left((\rho + \sigma(1-\rho))\,\pi_r + (1-\sigma)(1-\rho)\,\pi^*(x)\right) \\ 0 =& \sigma\left[w(x,\pi_r) + \beta_o V_1\left(\rho'\right) - \left(w(x,\pi^*(x)) + \beta_o V_1(0)\right)\right] \\ \rho' =& \frac{\rho}{\rho + (1-\rho)\,\sigma} \end{aligned}$$

which is the problem where the rule designer lacks commitment. The characterization of the equilibrium outcome follows from the characterization in provided in the proof of Proposition 3 for $\rho \in [0,\hat{\rho}] \cup \left[\rho_2^*, 1\right]$.

Part 3. Consider $\rho \in \left(\rho_3^*, \rho_2^*\right)$. In this region, we know that when rules are chosen without commitment, the rule designer chooses to pool if the residual horizon is $k = 1$ but it chooses to separate if the residual horizon is $k = 2$. This is because the optimizing type's incentive constraint is tighter when $k = 2$ than when $k = 1$ because $\Delta V_1(\rho) < \Delta V_0(\rho)$ as shown in Figure 5.

With commitment, the rule designer in period 0 can choose a stringent rule for period

1 to induce separation in period 1 and therefore relaxing the incentive constraint in period 0. The value associated with this plan is at least

$$\begin{aligned}\tilde{W}_2(\rho) &= w\left(x_1(\rho), \phi^{-1}(x_1(\rho))\right) + \beta W_0(\rho) + \beta^2 [\rho W_0(1) + (1-\rho) W_0(0)] \\ &> W_0(\rho) + \left(\beta + \beta^2\right) [\rho W_0(1) + (1-\rho) W_0(0)] \\ &= W_2(\rho)\end{aligned} \tag{69}$$

where the first inequality follows from the fact that for $\rho \in \left(\rho_3^*, \rho_2^*\right)$, the value of pooling in the twice repeated economy, $w\left(x_1(\rho), \phi^{-1}(x_1(\rho))\right) + \beta W_0(\rho)$, is higher than the value of separation, $W_0(\rho) + \beta [\rho W_0(1) + (1-\rho) W_0(0)]$, and the last line follows from the observation that without commitment it is optimal to separate in the first period so $W_2(\rho) = W_0(\rho) + \left(\beta + \beta^2\right) [\rho W_0(1) + (1-\rho) W_0(0)]$. This proves that the commitment solution does not coincide with the solution without commitment for $\rho \in \left(\rho_3^*, \rho_2^*\right)$. Under commitment, it is optimal to pool in the first period and promise to separate in the second period. Q.E.D.

The key insight of part 3 of the proposition is that the optimizing type's incentive constraint in period t is tighter if there is pooling in period $t+1$ as compared with the case in which there is separation in $t+1$ for sufficiently high levels of reputation. Thus the period t rule designer wants to have more stringent rules in period $t+1$ to induce separation. This channel does not operate in a two period economy because the terminal period's rule designer has no instruments to incentive the optimizing type to pool and thus there is always separation. Hence the first and second period rule designers agree and there is no time inconsistency problem.

F Signaling game and payoff types

In this section, we contrast our characterization of the optimal rule in Section 4 with two alternatives. First, we consider a signaling game in which the rule is chosen by the policy maker (which knows its type) instead of the rule designer, which is uncertain about the type of the policy maker. Second, we consider a model where the two types of policy makers differ in their preferences. In particular, policy makers can differ in their temptation to deviate ex-post because certain policy makers can better resist pressure from interest groups ex-post or have different preferences over outcomes than the social welfare function, as in the seminal Rogoff (1985) paper. We show that in both cases the equilibrium outcome has separation for all levels of reputation (under a reasonable refinement), in contrast with our main result that it is optimal to pool for low levels of reputation.

F.1 Comparison to a signaling game

We now consider a signaling game in which the rule is chosen by the policy maker (which knows its type) instead of the rule designer, which is uncertain about the type of the policy maker. If the rules are chosen by the policy makers, the commitment type (if sufficiently patient) will choose a rule that induces separation for *all* levels of reputation. In particular, it will prefer to separate for low levels of reputation even though the rule designer strictly prefers to pool. This result mirrors the one in Dovis and Kirpalani (2020a).

Proposition 10. *Under Assumptions 1 and 3, the outcome of the signaling game is such that the commitment and optimizing types follow different policies if either i) ρ is sufficiently high or ii) ρ is sufficiently small and $\beta_o = \beta_c \in (\underline{\beta}, \bar{\beta})$. Thus, in both cases there is separation after one period.*

The main idea here is that there are no dynamic gains for the commitment type of preserving uncertainty. The continuation value for the commitment type is always larger in a separating equilibrium as compared with pooling, as it can achieve the Ramsey outcome since the private agents know that they are facing the commitment type. However, there may still be static benefits of pooling when reputation is sufficiently low, as we saw in Section 4.1. But if the discount factor is sufficiently high ($\beta > \underline{\beta}$), the dynamic benefits outweigh the static losses. Note that for this to be an equilibrium we also need the optimizing type to strictly prefer to separate, which requires the discount factor to be low enough ($\beta < \overline{\beta}$). We show that $\underline{\beta} < \overline{\beta}$, since the optimizing type has additional static benefits of separating owing to the fact that it can choose its policy after the private agents have chosen their action.

Proof of Proposition 10. We consider the case with $w_x < 0$. Note that the statically optimal rule chosen by the commitment type is $\underline{\pi}$. To see why note that the first order condition for the commitment type is

$$\begin{aligned}
&w_\pi(x,\pi) + w_x(x,\pi)\frac{\phi'(\cdot)}{[1-\phi'(\cdot)(1-\rho)\pi^*_x(x)]} \\
\leqslant & w_\pi(x,\pi) + [\rho w_x(x,\pi) + (1-\rho)w_x(x,\pi^*(x))]\frac{\phi'(\cdot)}{[1-\phi'(\cdot)(1-\rho)\pi^*_x(x)]} \\
\leqslant & 0
\end{aligned}$$

where the first inequality follows from the assumption that $w_{x\pi} \geqslant 0$ and the last inequality from Assumption 2. Let

$$V^c_0(\rho) = w(x_0(\rho), \underline{\pi}),$$

be the value for the commitment type in the terminal period given the prior ρ where

$$x_0(\rho) = \phi\left(\rho\underline{\pi} + (1-\rho)\,\pi^*\left(x_0(\rho)\right)\right).$$

We can write the value for the commitment type if it chooses to separate as

$$V^c_{sep}(\rho) = w\left(x_0(\rho), \underline{\pi}\right) + \beta V^c_0(1),$$

while the value of pooling is

$$V^c_{pool}(\rho) = w\left(\phi\left(\pi_{ico}(\rho)\right), \pi_{ico}(\rho)\right) + \beta V^c_0(\rho)$$

where π_{ico} solves

$$w\left(\phi\left(\pi_{ico}(\rho)\right), \pi_{ico}(\rho)\right) + \beta V_0(\rho) = w\left(\phi\left(\pi_{ico}(\rho)\right), \pi^*\left(\phi\left(\pi_{ico}(\rho)\right)\right)\right) + \beta V_0(0).$$

Therefore,

$$V^c_{sep}(\rho) - V^c_{pool}(\rho) = \left[w\left(x_0(\rho), \underline{\pi}\right) - w\left(\phi\left(\pi_{ico}\right), \pi_{ico}\right)\right] + \beta\left[V^c_0(1) - V^c_0(\rho)\right]$$

First note that for ρ sufficiently large separating has both dynamic gains, $V^c_0(1) - V^c_0(\rho) > 0$, and static gains as $\left[w\left(x_0(\rho), \underline{\pi}\right) - w\left(\phi\left(\pi_{ico}\right), \pi_{ico}\right)\right] > 0$. In particular, for $\rho \to 1$ we have that the static gains of separating converge to

$$\left[w\left(\phi\left(\underline{\pi}\right), \underline{\pi}\right) - w\left(\phi\left(\pi_{ico}(1)\right), \pi_{ico}(1)\right)\right]$$

which is positive since under our assumption that the Ramsey outcome is not sustainable, $\pi_{ico}(\rho) < \underline{\pi}$. Consequently, for ρ large enough the commitment type will choose a stringent rule and thus there will be separation.

Next, given some ρ, the commitment type would like to separate if

$$\beta \geqslant \underline{\beta}(\rho) \equiv \frac{\left[w\left(\phi\left(\pi_{ico}(\rho)\right), \pi_{ico}(\rho)\right) - w\left(x_0(\rho), \underline{\pi}\right)\right]}{\left[V^c_0(1) - V^c_0(\rho)\right]}$$

To show that it is optimal for the optimizing type to separate at $\underline{\pi}$ it must be that

$$w\left(x_0(\rho), \pi^*\left(x_0(\rho)\right)\right) + \beta V_0(0) > w\left(x_0(\rho), \underline{\pi}\right) + \beta V_0(1)$$

(note that if the optimizing type mimics the commitment type the posterior jumps to one because we are constructing an equilibrium with separation) or

$$\beta < \bar{\beta}(\rho) \equiv \frac{[w(x_0(\rho), \pi^*(x_0(\rho))) - w(x_0(\rho), \underline{\pi})]}{V_0(1) - V_0(0)}$$

Therefore, the equilibrium outcome of the signaling game has separation if

$$\bar{\beta}(\rho) > \beta > \underline{\beta}(\rho)$$

Thus we need to show that such an interval exists. For $\rho \to 0$ we have

$$\lim_{\rho \to 0} \bar{\beta}(\rho) = \frac{w(x_0(0), \pi^*(x_0(0))) - w(x_0(0), \underline{\pi})}{V_0(1) - V_0(0)}$$
$$\lim_{\rho \to 0} \underline{\beta}(\rho) = \frac{w(x_0(0), \pi^*(x_0(0))) - w(x_0(0), \underline{\pi})}{V_0^c(1) - V_0^c(0)}$$

since $\pi_{ico}(\rho) \to \pi^*(x_0(0))$. Thus to compare $\bar{\beta}(0)$ and $\underline{\beta}(0)$ we only need to compare the denominators since the numerators are identical. In particular, $\bar{\beta}(0) > \underline{\beta}(0)$ if and only if $V_0(1) - V_0(0) < V_0^c(1) - V_0^c(0)$, or

$$w(x_0(1), \pi^*(x_0(1))) - w(x_0(0), \pi^*(x_0(0))) < w(x_0(1), \underline{\pi}) - w(x_0(0), \underline{\pi})$$

or

$$w(x_0(1), \pi^*(x_0(1))) - w(x_0(1), \underline{\pi}) < w(x_0(0), \pi^*(x_0(0))) - w(x_0(0), \underline{\pi})$$

Note that

$$w(x_0(0), \pi^*(x_0(0))) - w(x_0(0), \underline{\pi}) \geqslant w(x_0(0), \pi^*(x_0(1))) - w(x_0(0), \underline{\pi})$$

so we are left to show that

$$w(x_0(0), \pi^*(x_0(1))) - w(x_0(0), \underline{\pi}) > w(x_0(1), \pi^*(x_0(1))) - w(x_0(1), \underline{\pi})$$

Under Assumption 1, for $x_H > x_L$

$$\int_{[\underline{\pi}, \pi^*]} w_\pi(x_H, \pi)\, d\pi > \int_{[\underline{\pi}, \pi^*]} w_\pi(x_L, \pi)\, d\pi$$

Thus, since $x_0(0) > x_0(1)$ the inequality is satisfied. Q.E.D.

F.2 Payoff types

So far, we have modeled the commitment type as a policy maker that cannot deviate from the rule. An alternative to modeling the uncertainty about the policy maker's ability to follow the rule is to assume that the two types of policy makers differ in their preferences. In particular, the policy makers can differ in their temptation to deviate ex-post because certain policy makers can better resist pressure from interest groups ex-post or have different preferences over outcomes than the social welfare function, as in the seminal Rogoff (1985) paper.

We next show that with preference types and a reasonable equilibrium refinement, we have different outcomes than in our benchmark case. In particular, the equilibrium coincides with the outcome of the signaling game and there is separation for all levels of initial reputation.

We make our point in the context of the bailout example. Recall that the social welfare function is

$$w(x, \pi; \psi) = -\nu(x) + p(x) R_H - \psi(1 - p(x))(1 - \pi) - c(\pi),$$

where x is the banker's effort given by $\phi(\mathbb{E}\pi)$ for some ϕ with $\phi' < 0$, $\phi'' > 0$, $p(x)$ is the probability that the investment succeeds, and $\psi(1 - p(x))(1 - \pi)$ is the default cost that can be mitigated by transfers π. The parameter ψ controls the degree of time inconsistency: if $\psi = 0$ then the Ramsey outcome is sustainable because there are no benefits of deviating from the optimal plan ex-post. In contrast, if ψ is large then there is a much larger temptation to deviate ex-post.

Suppose now that there are two types of policy makers, each associated with a different value of ψ. The high cost type has $\psi = \psi_H > 0$, and the low cost type has $\psi = \psi_L = 0$. The low cost type then has no incentive to deviate ex-post and thus represents the commitment type in our baseline model. It also corresponds to the "conservative central banker" in Rogoff (1985) since if the private agents know they are facing the low cost type with probability one then the Ramsey outcome can be implemented. To keep the symmetry with the previous analyses, we assume that the social welfare function used by the rule designer to evaluate outcomes is $w(x, \pi; \psi_H)$.

Consider the twice-repeated problem. The characterization in the terminal period does not change relative to the case analyzed previously. Thus, the value for the rule designer is $W_0(\rho)$, where ρ is the prior of facing the low cost type, the value for the high cost type is $V_0(\rho; \psi_H) = V_0(\rho)$, and the value for the low cost type is $V_0(\rho; 0) = w(x_0(\rho), \pi = 0)$.

Consider now the rule designer's problem in the first period. The difference with problem (9) is that we have to add an incentive compatibility constraint for the low cost type,

$$w(x, \pi_r; 0) + \beta_o V_0 \left(\rho_c'; 0\right) \geqslant w(x, \pi; 0) + \beta_o V_0 \left(\rho'(\pi); 0\right) \quad \forall \pi,$$

where $\rho'(\pi)$ is the posterior after observing policy π and the low cost type's discount factor is β_o. Since $w_\pi(x, \pi; 0) = 0$ for all (x, π) then we can rewrite the constraint above as $\beta_o V_0(\rho_c'; 0) \geqslant \beta_o V_0(\rho'(\pi); 0)$ or, since $V_0(\rho; 0)$ is strictly increasing in ρ, as

$$\rho_c' = \rho'(\pi_r) \geqslant \rho'(\pi) \quad \forall \pi. \tag{70}$$

The incentive compatibility constraint for the low type, (70), is satisfied in the separation regime as $\rho_c' = 1$ so the rule designer can attain the same value. We now turn to analyze whether the pooling regime is feasible. The answer to this question depends on the specification of off-path beliefs. Clearly, it is possible to specify the off-path beliefs as follows

$$\rho'(\pi) = \begin{cases} \rho & \text{if } \pi = \pi_r \\ 0 & \text{if } \pi \neq \pi_r \end{cases}. \tag{71}$$

This choice is consistent with Bayes' rule on-path, trivially satisfies (70), and so supports the pooling outcome described above. An unappealing feature of (71) is that implementing more stringent policies ex-post reduces the policy maker's reputation. If we restrict to specifying beliefs such that $\rho'(\pi)$ is strictly decreasing in π then pooling is not feasible and the separating regime is the only solution for all levels of reputation. The restriction is intuitive as it imposes that if the deviation is relatively more advantageous for the low cost type then the posterior rises after enforcement