MAX WEBER PROGRAMME

EUI Working Papers

MWP 2010/16
MAX WEBER PROGRAMME

CREDIBILITY CONCERNS IN OPTIMAL POLICY DESIGN

Yang K. Lu

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE
MAX WEBER PROGRAMME

Credibilty Concerns in Optimal Policy Design

YANG K. LU

EUI Working Paper **MWP** 2010/16

ISSN 1830-7728

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu
cadmus.eui.eu

Abstract

This paper provides a transparent framework for thinking about optimal investment in reputation by a trustworthy government. This topic has been neglected by the existing literature on reputation games but is potentially a central issue in actual policymaking. A nonstandard feature of this model is that a "trustworthy" government (able to pre-commit) optimally chooses a policy announcement to which it commits. Thus, investing in reputation can be achieved by announcing policy plans that are costly and in turn are unlikely to be enacted if the government were an "opportunistic" type (unable to pre-commit). I identify the key trade-off in determining the optimal policy plans, which is mainly influenced by the difference in the time preference across types. When a trustworthy type is sufficiently more patient than an opportunistic type, there will be substantial reputation-building by the trustworthy type in the early stages of the game, which eventually leads to full separation from the opportunistic type.

Keywords

Imperfect credibility, reputation game, optimal taxation, time inconsistency.

JEL codes: E61, E62, D82

First version of this paper: November 2008
This version: May 2010

This paper is based on the first chapter of my doctoral thesis at Boston University. I owe a great debt of gratitude to Robert G. King for his constant guidance and support. The paper also benefits from discussions with Russell Cooper, Martin Cripps, Steve Davis, Piero Gottardi, Veronica Guerrieri, Peter Ireland, Bart Lipman, Lars Stole, and seminar participants at Boston University, Green Line Macro Meeting, UC-Santa Cruz, UChicago-Booth, FRB Richmond, FRB Kansas City, FRB Philidephia, Hong Kong University, Hong Kong University of Science and Technology, Chinese University of Hong Kong, City University of Hong Kong, European University Institute, Toulouse School of Economics and LUISS. All errors are my own.

Yang K.Lu
Max Weber Fellow, 2008-2009
Hong Kong University of Science and Technology

Introduction

When is it desirable for an optimizing trustworthy government to invest in reputation? This is an important question for actual policymakers, but game theory and macroeconomics have not yet provided an answer. Most of the literature on reputation games focuses on situations in which it is optimal for *opportunistic* types (types that have short-run incentives to break promises) to invest in reputation by forgoing short-run interests and mimicking trustworthy types (types that can commit to promises).[1] However, little has been said about the *trustworthy* type's optimal investment in reputation, since it is generally assumed to always play the same single action.

This paper provides an answer to the question by letting a trustworthy government optimally choose a policy announcement to which it commits.[2] I construct a simple taxation model in the tradition of past work on optimal policy with limited commitment by Barro and Gordon (1983), Fischer (1990), and Phelan (2006). Three important results follow.

First, given this additional policy instrument, I show that the trustworthy type, instead of choosing the same plan all the time, may strategically build up its reputation by announcing plans that would unlikely to be carried out if the opportunistic type were in place. Second, I characterize the conditions under which this strategic reputation-building occurs in equilibrium. Since committing to tough plans involves short-term costs, reputation-building by the trustworthy type is only optimal when it is sufficiently more patient with respect to tax revenues than its opportunistic counterpart. Third, I establish the uniqueness of the equilibrium announcement by adopting the refinement proposed by Mailath, Okuno-Fujiwara and Postelwaite (1993) as a way to resolve the well-known problem of equilibrium indeterminacy in signalling games.

The model is a finitely-repeated reputation game, where a government that seeks to maximize the present value of its tax revenues plays against a continuum of small private-sector households who decide whether or not to engage in production with a fixed cost each period. The government chooses a tax plan and announces it at the beginning of each period.[3] The households entertain the possibility that the government is one of two types, differing with regard to their ability to commit: a 'trustworthy' government, which by definition is sure to enact its announced plans; or an 'opportunistic' government, which retains the option of deviating from its announced plans. In the model, the optimal deviation is to the highest tax rate possible, representing confiscation by the government of all household production. The type of government is private information and is fixed throughout the whole game. Households update their beliefs in a Bayesian fashion after observing the government's action at the end of each period.

The government in the model uses the tax plan/announcement to control both the households' perception of the likelihood that the current plan will be carried out and how the households' beliefs concerning the government type will evolve once the plan is actually carried out. An important trade-off arises when the trustworthy government chooses its plans.[4] On the one hand, the trust-

[1] See for examples: Kreps and Wilson (1982), Milgrom and Roberts (1982), Barro (1986), Fudanburge and Levine (1992), Phelan (2006), Liu (2009), and Cripps, Mailath and Samulson (2004).

[2] This paper only considers pure strategies for policy announcements.

[3] Although within-period commitment is used here, the conjecture is that nothing will be changed if we assume commitment to a state-contingent plan for the entire future, given that there is neither uncertainty nor are there forward-looking constraints in the model. I thank Borys Grochulski for pointing this out.

[4] If the government were opportunistic, it would announce exactly the same planned tax rate as a trustworthy government would announce for two reasons. First, it is costless for an opportunistic government to imitate the announcements. Second, households know what a trustworthy government will announce so that any deviation from the anticipated rate will result in full revelation, which, in turn, will cost the opportunistic type the entire future stream of tax revenues. The detailed reasoning will become clear in the paper and this is where the refinement by

worthy government has an incentive to separate itself from the opportunistic type. In particular, it would like to set a *low* planned tax rate so that enacting the plan would be costly and thus unlikely if the opportunistic type were in place. This benefits the trustworthy type because by enacting the plan, it then conveys more information about its identity and improves its reputation. As a result, there will be more production and thus a greater flow of tax revenues in the future. On the other hand, the trustworthy government also has an incentive to pool with the opportunistic type. In particular, it would like to set a *high* planned tax rate, a plan that would be easily enacted even by the opportunistic type. This raises the credibility of the current plan and in turn reduces households' risk of confiscation. As a result, more households produce and more tax revenues are secured for the current period.

I analyze how this trade-off between future and current credibility is affected by the key parameters in the model, namely the patience levels of the two types of government, the elasticity of supply, the initial reputation of the government, and the time horizon of the game. Among these, I find that the relative patience of the trustworthy type with respect to that of the opportunistic type plays a key role in determining which side of the trade-off dominates.

When the trustworthy type is sufficiently more patient than the opportunistic type, the incentive to signal dominates, so that the Markov perfect equilibrium involves initial periods of rapid learning with low taxation. As a result of this aggressive reputation-building, the trustworthy type is able to sustain robust economic outcomes in the later stages of the game despite following a path of increasing planned tax rates, which converge to a constant when full credibility is achieved. By contrast, when the trustworthy and opportunistic types share a similar time preference, the incentive to pool dominates, so that the Markov perfect equilibrium involves little learning with high taxation throughout the whole game.[5]

This paper shares the view of Mailath and Samulson (2001) that reputation building is an exercise not only in pooling with good types but also in distinguishing oneself – separating – from bad types.[6] Focusing on reputation building as separation, Mailth and Samulson (2001) abstract from the strategic behaviour of the bad type, which is assumed to always play the same single action. My model, instead, allows both types to act strategically and thus explores the interaction between incentives to separate and incentives to pool. It is precisely the competition in reputation-building between trustworthy and opportunistic types that determines the optimal policies.

A few papers prior to this one also have both trustworthy and opportunistic types behaving strategically. Cukierman and Liviatan (1991) analyze the case where the optimal announcement by the trustworthy government reacts to its imperfect credibility. The current paper builds on their arguments by emphasizing an important reverse relationship: policy announcements can also be a powerful tool to manipulate government credibility. A similar idea was explored in my paper with King and Pasten (2008b), where we provided a numerical example in which reputation building by a trustworthy central bank implied gradualism in disinflation resembling the "Volker disinflation". The current paper, however, performs a more transparent and systematic analysis of the key trade-off facing the trustworthy type by using a sharply-focused example, and identifies the conditions under which reputation building by a trustworthy government occurs in equilibrium.[7]

Mailath et al. (1993) is employed.

[5] The equilibrium in this case is not exactly pooling for two reasons. First, the game has a finite horizon so that learning always occurs towards the end of the game when the reputation force becomes too weak to discipline the opportunistic type's behaviour. Second, if the opportunistic type is too impatient to care about reputation, it is unfeasible for the trustworthy type to force pooling even at the beginning of the game. However, in neither case is learning induced by the trustworthy type's active policymaking for the purpose of reputation building.

[6] See also Horner (2002).

[7] For example, King, Lu and Pasten (2008) use a random discount factor for the opportunistic central bank. But

Similarly to the current paper, in the literature on unsecured debt, Cole, Dow and English (1995), D'Erasmo (2007) and Chatterjee, Corbae and Rios-Rull (2008) have a difference in time preference across types as a central modelling component.[8] However, rather than focusing on the default or debt repayment decision, the current paper – in the context of unsecured debt – shows how the optimal level of debt can be determined by the borrower, who balances the trade-off between the current and future credibility of debt repayments.

In my model, the government has all the bargaining power with respect to the households, but it also has private information about its type. This feature links my analysis to the literature on the informed-principal problem pioneered by Myerson (1983) and Maskin and Tirole (1992). In common with this literature, the optimal contract is influenced by the principal's incentive to hide or reveal its private information, which thus leads to a pooling or separating equilibrium. The current paper, however, also includes mixed strategies by the opportunistic type as equilibria, so that it generates dynamics of credibility and provides a theory of reputation building.

The design of policy plans/announcements is the central focus of this paper, which relates it to models of strategic information transmission initiated by Crawford and Sobel (1982) and Sobel (1985). However, the announcement in my model is only cheap talk for the opportunistic type, but *not* for the trustworthy type because it is committed to act as it announces. This nonstandard feature of signalling differentiates the paper from the rest of the literature. In consequence, popular refinements such as the intuitive criterion lead to counterintuitive results.[9] Fortunately, a unique and plausible equilibrium can be reached when the refinement by Mailath, Okuno-Fujiwara and Postelwaite (1993) is applied.

It may seem easy to confuse this paper with models of repeated moral hazard with learning (Cukierman and Meltzer, 1986; Holmstrom, 1999) since in both cases reputation plays a major role in agents' decision-making. However, in Cukierman and Meltzer (1986), the government objective changes continuously over time and reputation is the speed with which the private sector detects the changes. By contrast, the government type in my model is fixed over time and reputation is defined as the perception of a government being trustworthy. This feature makes it possible for the opportunistic type to mimic and for pooling to be a potential equilibrium. With respect to the literature on career concern following Holmstrom (1999), in which information is imperfect but symmetric, the current paper addresses information *asymmetry* between governments and the private sector so that signalling becomes relevant in policymaking.

The remainder of the paper is organized as follows. Section 2 presents the setup of the game and defines the equilibrium concept. Section 3 uses a two-period example to illustrate the main results of the paper. Section 4 extends those results to a game with an arbitrary number of periods. Section 5 establishes the uniqueness of the equilibrium. Finally, Section 6 concludes.

A Reputation Game

In this section, I describe a fiscal policy game played between a government, which can be either trustworthy or opportunistic, and a competitive private sector composed of atomistic households. Unlike previous analyses of similar fiscal games (e.g. Phelan, 2006), where the trustworthy government simply sets an exogenously fixed tax rate, I work with a case in which the trustworthy

the mean discount factor has to be sufficiently low to get reputation building by the trustworthy central bank.

[8] See also Diamond (1989).

[9] For example, the intuitive criterion selects the worst equilibrium in which the government announces a 100 percent tax rate and no household produces; the refinement by Grossman and Perry (1986) rules out all sequential equilibria.

government is able to choose the optimal tax policy each period. One can then learnx how the tax path optimally responds to the dynamics of imperfect credibility. The equilibrium notion in the game will be Markov perfect.

Setup

There is a finite horizon economy, indexed by $t = 1, 2, ..., T$, in which a continuum of households play against a long-lived government that collects tax on household production. Within each period, the government moves first to announce a proportional tax plan $\tau_t \in [0, 1]$ on production for that period. The households observe the tax plan and decide whether to produce or not, according to their individual production costs. After the production decisions have been made, the government then levies the taxes on their outputs at rate $\tilde{\tau}_t$.

Households

Each household is indexed by i with total measure 1. For convenience, we can think of i as uniformly distributed over the interval [0,1]. The action profile for households is then a measurable function $a : [0, 1] \to \{0, 1\}$. Households only live for one period, but are able to observe the whole history of the game.[10] Each household is endowed with one indivisible unit of labor that they can put into production. The production technology is linear in labor with deterministic productivity normalized to 1. The production cost c_i for household i is a random draw each period from distribution $G(\cdot)$ with support $[0, 1]$.[11] It is convenient to work with the family of distributions $G(x) = x^{\gamma}$.[12] As discussed further below, the parameter γ is interpreted as the elasticity of supply.

If household i decides to produce, $a(i) = 1$, its payoff is the after-tax production net of the production cost, $r(\widetilde{\tau}, a(i) = 1) = (1 - \widetilde{\tau}) - c_i$. If household i does not produce, $a(i) = 0$, it gets zero, $r(\widetilde{\tau}, a(i) = 0) = 0$. We can use $\mu = \int_0^1 a(i)\, dG(c_i)$ to summarize the *aggregate participation of the households*. μ is publicly observable, but the individual production decision is private information to each household. Under this assumption, no individual household's action has an impact on another's decision, nor does it affect government action.

Government

The government lives for T periods and can be one of two types: trustworthy or opportunistic (TR and OP hereafter). Government type is private information and stays fixed throughout the game. The trustworthy government is assumed to always implement the announced tax rate τ_t at the end of each period, $\tilde{\tau}_t = \tau_t$. The opportunistic type, however, can either tax the production at rate τ_t, or confiscate all output, i.e. $\widetilde{\tau}_t = 1$, or even randomize between $\widetilde{\tau}_t = \tau_t$ and $\widetilde{\tau}_t = 1$. Let π_t denote *the probability of confiscation.*

The government makes decisions at two distinct points within a period. At the beginning of the period, it sets the tax plan τ_t for the current period. At the end of the period, it decides whether

[10] The assumption of short-lived households is not essential. As long as the rewards of production cannot be stored over periods, the maximization problem of households will be the same with a longer horizon.

[11] Macroeconomists generally study the effects of policies in settings where there is a smooth response on the part of the private sector to public interventions such as taxation and monetary policy. For example, the early 1980s analyses of monetary policy and reputation, such as Backus and Driffill (1985 a,b) and Barro (1986), worked in such settings. To this end, households in my model are heterogenous in their production costs.

[12] As will be seen later, this particular functional form of $G(\cdot)$ allows a close-form solution in this economy. At the same time, it abstracts from any variations in optimal policies that are not due to credibility concerns. Hence, we can focus solely on the interaction between optimal policy designs and credibility concerns.

the plan will be carried out, i.e. choosing π_t. The latter decision, however, is only relevant for the opportunistic type, since it is not an option for the trustworthy type by definition. Hence, we will refer to π_t as the action of the opportunistic type hereafter.

Both types of governments want to maximize their life-time tax revenues, discounted at rates β_{tr} and β_{op}, respectively. The objective is then defined as

$$\sum_{t=0}^{T} \beta_{type=tr,op}^{t} \; \widetilde{\tau}_t \mu_t. \tag{1}$$

Markov strategies and Markov perfect equilibria

Strategies are mappings from histories to actions. This paper restricts attention to Markov strategies which condition actions solely on payoff-relevant variables. In this incomplete information game, the only payoff-relevant variable is households' belief that the current government is of the trustworthy type.[13] Define this belief to be the state variable of the game: ρ. Under the Markov restriction, we can write the strategies of both the households and the government as functions of ρ: $\{\tau(\rho), \mu(\rho), \pi(\rho)\}$. Instead of specifying each household's strategy $a_i(\rho)$, we use the participation rate $\mu(\rho)$ as a convenient aggregation of individuals' strategies. This participation rate is also the tax base for the government, as suggested by (1) above.

Household beliefs and credibility concepts

New-born households at period t have the same information set as that of the previous generation at the end of period $t-1$. Thus their prior belief that the current government is trustworthy equals the state variable ρ_t. After observing the announced tax plan τ_t, households update this belief to $\rho_{t+0.5}(\tau_t)$. Then at the end of period t, after observing the tax action $\tilde{\tau}_t$, the belief is further updated to $\rho_{t+1}(\tilde{\tau}_t)$.

Why there can be belief updating right after the announcement is due to the fact that trustworthy and opportunistic types may potentially announce different tax rates. If so, the observed announcement will have information content about the current government type. However, to simplify matters at this stage of analysis, I impose a working assumption that *the opportunistic government always imitates the trustworthy type's announcement*. This assumption makes the tax announcement uninformative about the identity of the government so that $\rho_{t+0.5}(\tau_t) = \rho_t$ for any τ_t. Section 5 will revisit this assumption and argue that it is the equilibrium announcement strategy of the opportunistic type in an explicitly-modelled signalling game. Roughly speaking, the intuition behind this result is that the opportunistic type, with no cost of reneging on the announced rate, has an interest in keeping itself indistinguishable from the trustworthy type at the announcement stage. In addition, Section 5 will show that the equilibrium outcomes obtained under this assumption are the same as in the signalling game once the refinement of Mailath et. al. (1993) is adopted. Therefore, imposing this working assumption has no effect on either the equilibrium strategies or the equilibrium outcomes.

On the other hand, imposing this assumption significantly simplifies the analysis in two ways. First of all, since households' belief does not change after the tax announcement, their perception

[13] When households choose whether to produce or not, both the identity of the government and the announced tax rate are payoff-relevant variables. The identity of the government is not observable by households so it has to be replaced by households' best estimate of it. The announced tax rate itself is an optimal choice by the government, whose payoff depends on how many households produce and, in turn, their best estimate of the government type.

This extension of MPE to an incomplete information game is in the spirit of Ball (1995): "... actions depend on agents' best estimates of payoff-relevant variables".

of how likely the government will implement the announced tax rate, denoted by ψ_t, is simply:

$$\psi_t = \rho_t + (1 - \rho_t)(1 - \pi_t).$$

This is the sum of two elements: the likelihood of a trustworthy type of government being present so that τ_t will be implemented for sure, and the likelihood of an opportunistic type being present but mimicking with probability $(1 - \pi_t)$. Because ψ_t is the credibility of the current tax announcement, I call it "short-term credibility" so as to distinguish it from ρ_t, which measures the credibility of tax announcements in the long-run and is thus called "long-term credibility" in this paper.

The belief updating after observing the tax action can be characterized by Bayes' rule. In particular, if confiscation occurs at the end of the period, long-term credibility next period (ρ_{t+1}) will be zero, $\rho_{t+1} = 0$, since only the opportunistic type may confiscate. In the absence of confiscation, however, the evolution of long-term credibility will be:

$$\rho_{t+1} = \frac{\rho_t}{\rho_t + (1 - \rho_t)(1 - \pi_t)} = \frac{\rho_t}{\psi_t}, \tag{2}$$

where the marginal probability of observing $\widetilde{\tau}_t = \tau_t$ is short-term credibility by construction.

Second, the working assumption makes the trustworthy type a *leader* in the game, because the tax announcement, regardless of the true identity of the policy maker, is set to maximize the trustworthy government's payoff. The stage game played between the current government and the households can thus be analyzed as if there were three players in the game, with a trustworthy government setting the announcement first, and the households producing accordingly, followed by the opportunistic type deciding whether to confiscate, given the tax announcement and the households' production. [14]

Household strategies

As households make production decisions before the taxes are levied, they have to form expectations of the actual tax rate. I impose rational expectation and, in turn, symmetry across households, so the expected tax rate is the probability-weighted average between τ and 1, with the probability attached to τ being the government's short-term credibility ψ. The atomistic feature of each household eliminates any strategic or intertemporal concern in the production decision-making, because the decision does not individually affect the play of the government or the future value of its long-term credibility ρ. Therefore, the decision on the households' side is essentially static, with all households using a common threshold to decide whether to produce or not. Only those who receive cost c_i less than the expected after-tax output $(1 - \tau)\psi$ will produce, resulting in the optimized aggregate strategy:

$$\mu = G[(1 - \tau)\psi]. \tag{3}$$

With the distribution $G(x) = x^\gamma$, the participation rate is $[(1 - \tau)\psi]^\gamma$ so that the parameter γ is the elasticity of supply with respect to the expected reward to production, $[(1 - \tau)\psi]$.

[14] Even when the government in place is indeed trustworthy, it still needs to account for the reaction of the opportunistic type in terms of the confiscation probability. This is because the households, when making production decisions, are in general concerned about a possible confiscation by an opportunistic government.

Government strategies

In contrast to individual household, the government is both a big player, so that it is strategic, and a long-lived one, so that its optimization is intertemporal.

The opportunistic type obtains instant revenue gain, μ_t, by confiscating the current output. However, its true identity is then revealed and no production will occur in future periods.[15] The loss of future tax revenues can be avoided if the opportunistic government mimics the trustworthy type's behavior by taxing at the announced rate τ_t. This disciplined behavior increases its long-term credibility level to a new future state: $\rho_{t+1} = \rho_t/\psi_t$, generating a positive future value of tax revenues, denoted by V_{t+1}.[16] This trade-off between current and future tax revenues is summarized by what I call the "incentive compatibility constraints," which are stated as follows:

$$\begin{array}{rcll} \text{Pooling strategy (mimicking)} & \psi & = & 1 \text{ if } \mu \leqslant \tau\mu + \beta_{op}V'(\rho') \\ \text{Separating strategy} & \psi & = & \rho \text{ if } \mu \geqslant \tau\mu + \beta_{op}V'(\rho') \\ \text{Mixed strategy} & \psi & \in & (\rho, 1) \text{ if } \mu = \tau\mu + \beta_{op}V'(\rho'). \end{array} \tag{4}$$

Notice that the opportunistic type's strategy is stated in terms of ψ, instead of π, for two reasons. First, as will become clear later, using ψ turns out to be more convenient and intuitive. Second, there is a one-to-one mapping between π and ψ conditional on the long-term credibility level ρ. With $\psi = 1$, it mimics for sure: $\pi = 0$; and with $\psi = \rho$, it confiscates for sure: $\pi = 1$. Randomizing between mimicking and confiscating with $\pi \in (0, 1)$ corresponds to ψ lying within $(\rho, 1)$. Hence, we will use the strategy triple $\{\tau(\rho), \mu(\rho), \psi(\rho)\}$ to characterize the Markov perfect equilibrium (MPE) in the game, although we could have equivalently used $\{\tau(\rho), \mu(\rho), \pi(\rho)\}$.

The trustworthy type, as a leader in this game, is able to use its tax announcements τ to affect both household production μ and the opportunistic type's confiscation probability π, in order to maximize its life-time tax revenues. The impact of tax announcements on this objective works through various channels. Most obviously, the tax rate announced, if actually imposed, directly affects how much revenue a government can collect each period. This payoff relevance has a further impact on how likely the opportunistic type is to follow the announced rate. The likelihood of confiscation by the opportunistic type then determines its short-term credibility, so that both current output and the evolution of credibility are affected, as the households are rational and Bayesian. The trade-offs involved in determining the tax announcement, as well as the resulting equilibrium strategy, are the focus of Sections 3 and 4.

Markov perfect equilibrium

A Markov perfect equilibrium of the multistage game may be defined formally as the following.

Definition 1 *The sequence of Markov strategy triplets $\{\tau_t(\rho), \mu_t(\rho), \psi_t(\rho)\}_{t=1}^{T}$ is called Markov perfect equilibrium (MPE) if, at each stage of the game, $\tau_t(\rho)$ solves the trustworthy government's optimization:*

$$W_t(\rho) = \max_{\tau_t} \{\tau_t\mu_t + \beta_{tr}W_{t+1}(\rho/\psi_t)\},$$

[15] If the government is known for sure to be of the opportunistic type $\rho = 0$, the uniqe MPE in this game is zero production. To see that, consider any period with $\tau < 1$. If $\pi = 1$, household optimization implies $\mu = 0$. If $\pi < 1$, government optimization requires $\tau\mu + \beta_{op}V(0) >= \mu + \beta_{op}V(0)$, which only holds for $\mu = 0$. Given $\mu = 0$, the weak type is indifferent among any choice of π. Letting $\pi = 1$, it justifies the zero production choice of the households. If no one produces in this game, the life-time tax revenue will be zero: $V(0) = 0$ for any period.

[16] The subindex of the value function indicates that its functional form is time-varing due to the finite horizon of the game.

with $\{\mu_t(\rho), \psi_t(\rho)\}$ *satisfying the incentive compatibility constraints of both households and the opportunistic government, (3) and (4), subject to the evolution of credibility (2).*

Backward induction can be applied to obtain a MPE solution.[17] In each period, the trustworthy government takes the future value function as given when setting the optimal tax plan to balance its current and future payoffs. The current tax plan affects the equilibrium reactions of both the opportunistic type and the households through conditions (3) and (4), which in turn determines the current tax revenues $\tau\mu$. This momentary payoff is weighted against the future value of revenues, which is also affected by the current tax plan through its impact on the evolution of credibility governed by Bayes' rule (2).

A Two-Period Model

This section solves a two-period version of the model to illustrate the basic trade-offs facing a trustworthy government with imperfect credibility. On the one hand, the tax plan tends to be higher than a benchmark rate set by a government with no credibility concerns, because implementing a plan in which households have little faith would be costly. On the other hand, a tax plan lower than the benchmark rate, once implemented, can be a powerful tool to signal the government's type and thus changes the evolution of households belief. The relative strength of these offsetting forces critically depends on the difference between time preferences of the two types β_{tr}/β_{op}, the elasticity of supply with respect to tax changes γ, and the long-term credibility level ρ. The conditions under which it is optimal for the trustworthy type to either accommodate imperfect credibility or invest in future reputation will be characterized.

Optimal tax plan with full credibility

First consider a benchmark case where a trustworthy government possesses full credibility, $\rho = \psi = 1$. The expected after-tax output is $(1-\tau)$ in this case, which determines the fraction of households who produce, according to the optimized aggregate strategy (3). The production base thus decreases with the announced tax rate, as shown in Panel A of Figure 1. This feature replicates the celebrated Laffer curve, which traces an inverted-U relationship between taxes and revenues. Panel B in Figure 1 depicts one example of this Laffer curve. The optimal tax rate maximizes revenues and is thus:

$$\tau^*_{UC} = \arg\max_{\tau} \tau G(1-\tau).$$

I label such a rate "unconstrained optimal" because it is free of credibility concerns. If we use the specific cost distribution $G(x) = x^\gamma$, $\tau^*_{UC} = (\gamma+1)^{-1}$.

[17] The government in my model announces a tax plan at the beginning of each period rather than a path of tax plans for the entire future at time 0 (although these two cases are conjectured to be equivalent given there are neither shocks nor forward-looking constraints in my model.) The former case can be thought of as a limited commitment case in which the trustworthy type has no access to the intertemporal commitment technology, even though it can commit to the current period's tax plan. Thus, when the trustworthy government sets a value τ_t in period t, it recognizes that $\{\tau^*_s\}^T_{s>t}$ will be chosen by future trustworthy governments, with similar objectives. Therefore, backward induction can be applied.

Optimal tax plan in the final period

Let us start with the final period, $t = 2$. With some inherited credibility, the opportunistic type will always confiscate, irrespective of ρ and τ. Given that behavior, its short-term credibility is fixed at the level ρ. The homogeneity feature of $G(\cdot)$ ensures that the optimal tax plan is always equal to the unconstrained optimal rate for any fixed short-term credibility level: $\tau^*_{UC} = \arg\max \tau G\left[(1-\tau)\,\bar{\psi}\right]$. Thus, the trustworthy type sets $\tau^*_2 = \tau^*_{UC}$ to maximize current tax revenues, which yields the value functions of both types as:

$$\begin{aligned} V_2(\rho_2) &= (1-\tau^*_{UC})^{\gamma}\rho_2^{\gamma} \\ W_2(\rho_2) &= \tau^*_{UC}V_2(\rho_2). \end{aligned} \tag{5}$$

This property, arising from the homogeneity of $G(\cdot)$, is desirable since any variation in optimal tax policies obtained later can be attributed to the endogeneity of short-term credibility. Endogenous short-term credibility captures the opportunistic type's reaction to the tax plan and is the essence of what I call "credibility concerns." In consequence, the deviation of the optimal tax rate from τ^*_{UC} provides a measure of credibility concerns.

Optimal tax plan with credibility concerns

Moving back to the first period, $t = 1$, the trustworthy government takes $W_2(\rho)$ and $V_2(\rho)$ as given, and chooses τ_1 by solving:

$$\max_{\tau_1} \tau_1\mu_1 + \beta_{tr}W_2(\rho_1/\psi_1).$$

The opportunistic type, in contrast to the final period, no longer has a strictly dominant strategy, since its next-period payoff critically depends on its tax action in the current period. As a result, the likelihood (including 0 and 1) that the opportunistic type will confiscate varies with different levels of long-term credibility ρ_1 and tax plan τ_1. The decision of the opportunistic type in turn affects the expected tax rate, which determines how many households will produce in that period. Therefore, when the trustworthy type decides the optimal tax plan, it has to take into account the best reaction of the opportunistic type, i.e. its incentive compatibility constraint, which takes one of the following three forms:

1) *a mixed outcome* $\psi_1 \in (\rho, 1)$ is induced by a planned tax rate that makes the opportunistic type indifferent:

$$\tau_1\mu_1 + \beta_{op}V_2(\rho_1/\psi_1) = \mu_1$$

$$\text{with } \mu_1 = G\left[(1-\tau_1)\,\psi_1\right] \text{ and } \psi_1(\tau_1, \rho_1) \in (\rho_1, 1).$$

2) *a pooling outcome* $\psi_1 = 1$ is induced by a planned tax rate that is bounded below by:

$$\tau_1\mu_1 + \beta_{op}V_2(\rho_1) \geqslant \mu_1 \text{ with } \mu_1 = G[1-\tau_1].$$

3) *a separating outcome* $\psi_1 = \rho_1$ is induced by a planned tax rate that is bounded above by:

$$\tau_1\mu_1 + \beta_{op}V_2(1) \leqslant \mu_1 \text{ with } \mu_1 = G\left[(1-\tau_1)\,\rho_1\right].$$

The trustworthy type first maximizes for each of the three outcomes, and then picks the optimal $\tau_1(\rho_1)$ that generates the highest value among the three cases.

If τ^*_{UC} satisfies either the pooling or the separating incentive compatibility constraint, i.e. $\beta_{op}\rho_1^{\gamma} \geqslant (1-\tau^*_{UC})$ or $\beta_{op} \leqslant (1-\tau^*_{UC})\,\rho_1^{\gamma}$, τ^*_{UC} is the optimal rate among those that generate pooling or separating outcomes. To see this, notice that when ψ_1 is fixed, the trustworthy government's tax revenues, $\tau_1(1-\tau_1)^\gamma(\bar{\psi}_1)^\gamma$, are maximized at τ_{UC}^* and its continuation value, $\beta_{tr}W_2(\rho_1/\bar{\psi}_1)$, is independent of the current tax plan. This solution to the unconstrained maximization does not change after imposing the incentive compatibility constraint because the constraint is not binding at τ_{UC}^*.

However, if τ_{UC}^* fails to support either a pooling or a separating outcome, that is when

$$\rho_1 \leqslant \min\left\{\left[\beta_{op}/(1-\tau_{UC}^*)\right]^{1/\gamma}, \left[(1-\tau_{UC}^*)/\beta_{op}\right]^{1/\gamma}\right\}, \tag{6}$$

the optimal tax plan will make the incentive compatibility constraint binding, i.e. the opportunistic type will be indifferent between mimicking and confiscating. The following lemma formalizes this property.

Lemma 1 *When τ_{UC}^* does not satisfy the incentive compatibility constraint of either a pooling or a separating outcome, the problem of the trustworthy type can be simplified to:*

$$\max_{\tau_1\in[0,1]} \tau_1\mu_1 + \beta_{tr}W_2(\rho_1/\psi_1) \tag{7}$$

$$s.t.\ \tau_1\mu_1 + \beta_{op}V_2(\rho_1/\psi_1) = \mu_1 \text{ with } \mu_1 = (1-\tau_1)^\gamma(\psi_1)^\gamma \text{ and } \psi_1 \in [\rho_1, 1]. \tag{8}$$

Proof. As shown above, when short-term credibility ψ_1 is fixed, the optimal solution to the unconstrained maximization (7) is τ_{UC}^*. In the case of having a pooling or a separating outcome, imposing the incentive compatibility constraint will alter the optimal solution if τ_{UC}^* makes the constraint binding. Therefore, maximization dictates that the trustworthy type makes the least deviation possible from τ_{UC}^* to meet the pooling or separating incentive compatibility constraints. That is, $\hat{\tau}_1$ such that $\hat{\tau}_1\mu_1(\hat{\tau}_1,\rho_1) + \beta_{op}V_2(1) = \mu_1(\hat{\tau}_1,\rho_1)$ is the optimal rate among those that generate separating outcomes, and $\tilde{\tau}_1$ such that $\tilde{\tau}_1\mu_1(\tilde{\tau}_1,1) + \beta_{op}V_2(\rho_1) = \mu_1(\tilde{\tau}_1,1)$ is the optimal rate among those that generate pooling outcomes. In the case of a mixed outcome, the opportunistic type is already indifferent between mimicking and confiscating. Therefore, in any case, the indifference condition for the opportunistic type is satisfied. ∎

The trade-offs in setting the optimal tax plan

In the more interesting case when the optimal tax plan is different from the unconstrained optimal rate, the lemma says that the trustworthy type always sets the tax plan such that the opportunistic type is indifferent between mimicking and confiscating. The indifference condition (8) thus determines the best reaction of the opportunistic type to the tax plan τ_1 given long-term credibility ρ_1, which can be captured by the endogenous variation of short-term credibility ψ_1 as a function of both τ_1 and ρ_1 :

$$\psi_1(\tau_1,\rho_1) = \left(\frac{\beta_{op}(1-\tau_{UC}^*)^\gamma}{(1-\tau_1)^{\gamma+1}}\right)^{\frac{1}{2\gamma}}\rho_1^{\frac{1}{2}}. \tag{9}$$

It is easy to see that ψ_1 increases with τ_1:

$$\frac{\partial\psi_1}{\partial\tau_1} \geqslant 0.$$

The intuition is the following. Holding ψ_1 constant, decreasing the tax plan τ_1 increases the tax-revenue gap between following and deviating from the plan, and in turn raises the temptation to deviate. It thus increases the probability of confiscation π_1 and implies lower short-term credibility ψ_1 to maintain the indifference.

Lower short-term credibility ψ_1 in the current tax plan increases the households' risk of confiscation after production and thus depresses the tax base $\mu_1 = (1-\tau_1)^{\gamma}\psi_1^{\gamma}$. Therefore, while a lower tax plan makes it more rewarding to produce, it is less likely to be carried out. This credibility effect offsets the conventional tax effect on the production base as described by the Laffer curve. As a result, the marginal gain from raising the planned tax rate is larger when short-term credibility responds to it endogenously than when short-term credibility is fixed as in the full credibility case:[18]

$$\frac{\partial \tau_1 \mu_1}{\partial \tau_1} = \frac{\partial \tau_1 G\left[(1-\tau_1)\psi_1\right]}{\partial \tau_1} \geqslant \frac{\partial \tau_1 G\left[(1-\tau_1)\right]}{\partial \tau_1} \geqslant 0.$$

The tax plan maximizing current tax revenues is therefore higher than the unconstrained optimal rate τ_{UC}^*, with the discrepancy measuring the extent to which a trustworthy government accommodates imperfect credibility.

Announcing a higher tax plan may benefit current tax revenues, but it impedes the transmission of information from the government to households. Because a higher tax plan is more likely to be carried out by the opportunistic type as well, implementing it does not convey much information about the government's type. The Bayesian learning equation (2) reflects such a learning effect, where higher short-term credibility ψ_1 reduces future long-term credibility ρ_2. Since the trustworthy type's tax revenues in the final period increase with its long-term credibility, a higher tax plan today imposes a loss in terms of future revenues through slower growth in its long-term credibility:

$$\frac{\partial \beta_{tr} W_2\left(\rho_1/\psi_1\right)}{\partial \tau_1} \leqslant 0.$$

Therefore, the credibility and learning effect of a tax plan, both of which stem from the optimal reaction of the opportunistic type, constitute the trade-offs facing the trustworthy government. In determining the tax plan, it has to weigh current revenues against future continuation values.

The key determinants in the trade-offs

What is the overall effect of these trade-offs on the optimal tax plan? Is it optimal for the trustworthy type to lower the tax plan relative to the unconstrained optimal rate τ_{UC}^*, as a way of investing in future reputation? Or is it optimal to raise the tax plan relative to τ_{UC}^* so as to secure more current tax revenues? I will argue that the two most important parameters that determine the overall effect are: 1) the elasticity of supply γ relative to the number of future periods (set to 1 in the two-period case); and 2) the ratio of time discount factors β_{tr}/β_{op}.

To see this, it is useful to observe that the trustworthy type's payoff differs from the opportunistic type's value of mimicking only in its weight on future continuation values:

$$W_1(\rho) = \max_{\tau_1}\left\{\tau_1\mu_1 + \left[\frac{\beta_{tr}}{\beta_{op}}\tau_{UC}^*\right]\beta_{op}V_2\left(\rho_1/\psi_1\right)\right\}.$$

Furthermore, because the opportunistic type is indifferent, the value of mimicking equals the current

[18] The derivatives are positive because a reasonable tax rate should always be on the increasing side of the Laffer curve.

production base, so does its reaction to tax changes:

$$\frac{\partial\left[\tau_1\mu_1+\beta_{op}V_2\left(\rho_1/\psi_1\right)\right]}{\partial\tau_1}=\frac{\partial\mu_1}{\partial\tau_1}.$$

Therefore, if $\left(\beta_{tr}/\beta_{op}\right)\tau^*_{UC}=1$, the payoff to the trustworthy type is equal to the current production base, whose reaction to tax changes is determined by the elasticity of supply γ. On the other hand, given the elasticity of supply γ, the larger the ratio of time preference β_{tr}/β_{op}, the greater the weight that the trustworthy type assigns to its future continuation value, which in turn enhances its incentive to invest in future reputation by lowering the planned tax rate today. β_{tr}/β_{op} is thus a measure of the trustworthy type's incentive to signal.

Recall that the reaction of the production base to tax changes is a compound result of two competing forces: the credibility effect, working through endogenous short-term credibility, and the standard Laffer-curve effect, working through households' after-tax payoff. When the supply is elastic, $\gamma>1$, the standard Laffer-curve effect dominates and the production base reacts negatively to a tax increase, $\partial\mu_1/\partial\tau_1<0$. However, when the supply is relatively rigid, $\gamma\leqslant 1$, the credibility effect takes over and overturns the conventional reaction of the production base to a tax increase. It now increases with the planned tax rate: $\partial\mu_1/\partial\tau_1\geqslant 0$.

When $\partial\mu_1/\partial\tau_1<0$, decreasing the planned tax rate relative to τ^*_{UC} increases μ_1, and in turn increases the payoff to the trustworthy type if $\left(\beta_{tr}/\beta_{op}\right)\tau^*_{UC}=1$. In this case, the optimal tax plan is lower than τ^*_{UC} as long as the incentive for the trustworthy type to signal is not too weak, i.e. β_{tr}/β_{op} is not too small compared to $(\tau^*_{UC})^{-1}$. When $\partial\mu_1/\partial\tau_1\geqslant 0$, however, decreasing the planned tax rate reduces μ_1, and in turn reduces the payoff to the trustworthy type if $\left(\beta_{tr}/\beta_{op}\right)\tau^*_{UC}=1$. In this case, the optimal planned tax rate is higher than τ^*_{UC} as long as the incentive for the trustworthy type to signal is not too strong, i.e. β_{tr}/β_{op} is not too large compared to $(\tau^*_{UC})^{-1}$.

As the nature of the trustworthy type's optimization differentiates between low and high elasticity of supply, I will treat the two cases separately in discussing the Markov Perfect Equilibrium.

MPE when $\gamma>1$

First, take the case when $\gamma>1$ so that $\partial\mu_1/\partial\tau_1<0$. In this case, there will be three ranges of the state variable ρ_1, labelled as "credibility regions", each with its own form of equilibrium. The property of the MPE for each region depends on the ratio of time discount factors, as summarized by the following proposition. The proof is constructive.

Proposition 1 *Under the assumption $\gamma>1$, there exists a unique MPE in period 1 with the following properties:*

i. If $\beta_{tr}/\beta_{op}\in\left(0,\left(\gamma+1\right)/\gamma\right)$ and $\beta_{op}\geqslant 1-\tau^*_{UC}$, $\tau^*_1\geqslant\tau^*_{UC}$ and the MPE is:

"mixed":	$\tau^*_1\left(0\leqslant\rho_1\leqslant l_1\right)=\bar{\tau}_1;\ \psi^*_1=\psi_1(\bar{\tau}_1,\rho_1)$
"constrained pooling":	$\tau^*_1\left(l_1\leqslant\rho_1\leqslant h_1\right)=\tilde{\tau}_1\left(\rho_1\right);\ \psi^*_1=1$
"unconstrained pooling":	$\tau^*_1\left(h_1\leqslant\rho_1\leqslant 1\right)=\tau^*_{UC};\ \psi^*_1=1$

ii. If $\beta_{tr}/\beta_{op}\in\left[\left(\gamma+1\right)/\gamma,\gamma+1\right]$, $\tau^*_1\leqslant\tau^*_{UC}$ and the MPE is:

"mixed":	$\tau^*_1\left(0\leqslant\rho_1\leqslant l_1^{-1}\right)=\bar{\tau}_1;\ \psi^*_1=\psi_1\left(\bar{\tau}_1,\rho_1\right)$
"constrained separating":	$\tau^*_1\left(l_1^{-1}<\rho_1<h_1^{-1}\right)=\hat{\tau}_1\left(\rho_1\right);\ \psi^*_1=\rho_1$
"unconstrained separating":	$\tau^*_1\left(h_1^{-1}\leqslant\rho_1\leqslant 1\right)=\tau^*_{UC};\ \psi^*_1=\rho_1$

iii. If $\beta_{tr}/\beta_{op} \geqslant \gamma+1$, the MPE is the same as case ii, except that there is no mixed MPE any more.

Household production is $\mu_1^* = [(1-\tau_1^*)\,\psi_1^*]^\gamma$ in all cases and

$$\bar{\tau}_1 = \frac{2-\beta_{tr}/\beta_{op}}{\gamma+1-\beta_{tr}/\beta_{op}}; \tag{10}$$

$$\tilde{\tau}_1(\rho_1) = 1 - \left[\beta_{op}(1-\tau_{UC}^*)^\gamma \rho_1^\gamma\right]^{\frac{1}{\gamma+1}} \text{ and } \hat{\tau}_1(\rho_1) = 1 - \left[\beta_{op}(1-\tau_{UC}^*)^\gamma \rho_1^{-\gamma}\right]^{\frac{1}{\gamma+1}}; \tag{11}$$

$$h_1 = \left[\frac{1-\tau_{UC}^*}{\beta_{op}}\right]^{1/\gamma} \text{ and } l_1 = \left[\frac{(1-\bar{\tau}_1)^{\gamma+1}}{\beta_{op}\left(1-\tau_{UC}^*\right)^\gamma}\right]^{1/\gamma} \tag{12}$$

Proof. When the unconstrained optimal rate τ_{UC}^* does not support either a separating or a pooling outcome, i.e. $\rho_1 \leqslant \min\left\{h_1, h_1^{-1}\right\}$, the optimal tax rate solving the problem of the trustworthy type as described in (7) is $\bar{\tau}_1$ if we ignore the boundary condition $\psi_1 \in [\rho_1, 1]$. This rate decreases with both γ and β_{tr}/β_{op}. It is higher than τ_{UC}^* when $\beta_{tr}/\beta_{op} \leqslant (\gamma+1)/\gamma$.

Imposing the boundary condition, $\bar{\tau}_1$ is thus feasible as long as it implies short-term credibility $\psi_1(\bar{\tau}_1, \rho_1)$ within $[\rho_1, 1]$, i.e. $\rho_1 \leqslant \min\left\{l_1, l_1^{-1}\right\}$. Therefore, *the MPE is mixed with planned tax rate $\bar{\tau}_1$ when $\rho_1 \leqslant \min\left\{l_1, l_1^{-1}\right\}$.* l_1 decreases with β_{op} but increases with β_{tr}/β_{op}.

Whenever the boundary condition of $\psi_1(\bar{\tau}_1, \rho_1)$ binds, the MPE is either pooling or separating, depending on which side of the boundary is binding. If $\psi_1(\bar{\tau}_1, \rho_1) < \rho_1$, i.e. $\rho_1 > l_1^{-1}$, a separating outcome is the best for the trustworthy type. If $\psi_1(\bar{\tau}_1, \rho_1) > 1$, i.e. $\rho_1 > l_1$, a pooling outcome is the best. The optimal tax plans for generating separating and pooling outcomes are the least deviations from the unconstrained optimal rate, i.e. $\hat{\tau}_1$ and $\tilde{\tau}_1$, as proved in Lemma 1. $\hat{\tau}_1$ increases with ρ_1 and $\tilde{\tau}_1$ decreases with ρ_1.

Now we can combine the results established above to prove the proposition.

$l_1 \leqslant 1$ and $h_1 \leqslant 1$ if $\beta_{tr}/\beta_{op} \leqslant (\gamma+1)/\gamma$ and $\beta_{op} \geqslant 1-\tau_{UC}^*$. In this case, there exists a credibility region $\rho_1 > l_1$ such that the MPE is pooling. Within this region, when long-term credibility is high enough – $\rho_1 > h_1$, τ_{UC}^* can induce a pooling outcome and is thus the optimal rate; when $\rho_1 \in [l_1, h_1]$, τ_{UC}^* does not support a pooling outcome and thus $\tilde{\tau}_1$ is the optimal rate.

$l_1 \geqslant 1$ and $h_1 \geqslant 1$ if $\beta_{tr}/\beta_{op} \geqslant (\gamma+1)/\gamma$. In this case, β_{op} is smaller than $1-\tau_{UC}^*$ since β_{tr} cannot be greater than 1. There thus exists a credibility region $\rho_1 > l_1^{-1}$ such that the MPE is separating. Within this region, the separating outcome is constrained and induced by $\hat{\tau}_1$ as long as $\rho_1 < h_1^{-1}$. Once long-term credibility exceeds h_1^{-1}, τ_{UC}^* becomes the optimal rate.

When $\beta_{tr}/\beta_{op} \geqslant \gamma+1$, $\left(\beta_{tr}/\beta_{op}\right)\tau_{UC}^*$ is greater than 1. The trustworthy type thus puts more weight on future continuation values than the opportunistic type. Because decreasing the tax plan always raises the value of mimicking when $\gamma > 1$, it certainly benefits the trustworthy type. In addition, $\beta_{tr}/\beta_{op} \geqslant \gamma+1$ implies $\beta_{op} < 1-\tau_{UC}^*$. So in the region where an unconstrained outcome is feasible, the MPE is also separating. Therefore, with any level of long-term credibility, having a separating outcome is the dominant choice for the trustworthy type. ■

Figure 2 plots two representative policy functions $\tau_1^*(\rho_1)$ associated with cases i and ii. The left-hand panel is the case of equally patient government types: $\beta_{tr} = \beta_{op} > 1-\tau_{UC}^*$. The right-hand panel is the case where the trustworthy type is sufficiently more patient than the opportunistic type: $\beta_{tr}/\beta_{op} > (\gamma+1)\gamma^{-1}$. The elasticity of supply is $\gamma = 2$ so that the unconstrained optimal tax rate τ_{UC}^* is 1/3.

The most notable difference between these two functions is their opposite monotonicity. With equally patient government types, the optimal tax plan decreases as credibility increases, until it reaches the unconstrained optimal rate τ^*_{UC}. However, with the opportunistic type less patient than the trustworthy one, the optimal tax plan increases with credibility until it reaches the unconstrained optimal rate τ^*_{UC}.

This contrasting feature again stems from the difference in the extent to which the trustworthy type is willing to invest in credibility capital. When it is infeasible to build up reputation rapidly enough to compensate for the current revenue loss, the trustworthy type will use the tax plan to sustain high current revenues. In turn, the optimal tax plan needs to be higher than τ^*_{UC} to prevent the opportunistic type from confiscating outputs so that short-term credibility is large enough to induce a decent participation rate. This is the case when $\beta_{tr} = \beta_{op}$. By contrast, when $\beta_{tr} > \beta_{op}$, the main goal of the tax plan is to increase the probability of confiscation by the opportunistic type so as to stimulate rapid growth in long-term credibility. As a result, the optimal tax plan is lower than τ^*_{UC} in order to impose a higher cost on the opportunistic government if it were to mimic the trustworthy type's tax action.

This divergence in the purpose of the optimal tax plans naturally results in the opposite patterns of monotonicity. In both cases, the gap between the optimal tax plan and the unconstrained optimal rate τ^*_{UC} shrinks as long-term credibility grows. This reflects the diminishing credibility concerns of the trustworthy type in designing an optimal tax plan, as its credibility capital accumulates.

What remains to be discussed is the case when $\beta_{tr}/\beta_{op} < (\gamma+1)\,\gamma^{-1}$ and $\beta_{op} < 1 - \tau^*_{UC}$. By and large, the overall pattern of the policy function still follows the left-hand panel in Figure 2 (Case i). That is, the optimal tax plan decreases with long-term credibility, starting at rate $\bar{\tau}_1$ when ρ_1 is low and ending at the unconstrained optimal rate τ^*_{UC} when ρ_1 is high. But there are two major differences in equilibrium outcomes when compared to Case i. First, it is now a *separating* outcome that is induced by τ^*_{UC}, as a result of the relative impatience of the opportunistic type. Second, constrained pooling outcomes can only occur in the middle range of long term credibility if $l_1 \leqslant 1$, i.e. if β_{op} is relatively high or β_{tr}/β_{op} is relatively low.[19] This is because pooling outcomes are only desirable for the trustworthy type when the incentive to mimic (measured by β_{op}) is relatively strong compared to the incentive to signal (measured by β_{tr}/β_{op}). If this is the case, the trustworthy type's choice of a pooling or a separating outcome will depend not only on the ratio of time preference β_{tr}/β_{op}, but also on long-term credibility ρ_1.

MPE when $\gamma \leqslant 1$

We now turn to the case when household production is more rigid, $\gamma \leqslant 1$ so that $\partial\mu_1/\partial\tau_1 \geqslant 0$. This is an important case to study since $\gamma \leqslant n$, which implies $\partial\mu_n/\partial\tau_n \geqslant 0$ with n denoting the number of future periods, will be a common situation in a multi-period model. In addition, long-term credibility plays a bigger role in determining the property of a MPE in this case compared to the one with $\gamma > 1$. The following proposition characterizes the MPE. The proof is constructive.

[19] If $l_1 \leqslant 1$, we have to compare the payoff to the trustworthy type associated with the unconstrained separating outcomes to both the mixed outcomes and the constrained pooling outcomes. This is because the region where the constrained pooling outcomes may be optimal is always a subset of the region where τ^*_{UC} can induce a separating outcome. To see this, start with $\beta_{tr}/\beta_{op} = (\gamma+1)/\gamma$, the highest possible level in the current case. The constrained pooling region then coincides with the unconstrained separating region. Now, if we decrease β_{tr}/β_{op}, a lower β_{op} is needed to keep l_1 constant. Hence, with the constrained pooling region fixed at $[l_1, 1]$, decreasing β_{tr}/β_{op} expands the unconstrained separating region. It can be shown that the unconstrained separating outcomes only dominate when the long-term credibility is above a cutoff $\bar{\rho}_1$, where $\bar{\rho}_1 > l_1$.

If $l_1 > 1$, a constrained separating region exists but it will be a subset of the unconstrained separating region. Because it is always optimal to use τ^*_{UC} to induce a separating outcome whenever possible, the MPE in this case will not include the constrained separating outcomes.

Proposition 2 *Under the assumption $\gamma \leqslant 1$, there exists a unique MPE in period 1 with the following properties:*

i. With low $\rho_1 \leqslant L_0$, the MPE is "constrained pooling" $\{\tau_1^* = \tilde{\tau}_1(\rho_1), \psi_1^* = 1\}$.

ii. With moderate $\rho_1 \in [L_0, H_1]$, the MPE is "constrained pooling" except in the following cases:

a) if $\beta_{tr}/\beta_{op} \geqslant A_1(\beta_{op}, \gamma)$, the MPE is

$$\begin{array}{ll} \text{"mixed":} & \tau_1^*(L_0 \leqslant \rho_1 \leqslant L_1) = 0;\ \psi_1^* = \psi_1(0, \rho_1); \\ \text{"constrained separating":} & \tau_1^*(L_1 \leqslant \rho_1 \leqslant M_1) = \hat{\tau}_1(\rho_1);\ \psi_1^* = \rho_1. \end{array}$$

b) if $\beta_{tr}/\beta_{op} \geqslant A_2(\beta_{op}, \gamma) \geqslant A_1(\beta_{op}, \gamma)$, the MPE is

$$\text{"constrained separating":} \quad \tau_1^*(M_1 \leqslant \rho_1 \leqslant H_1) = \hat{\tau}_1(\rho_1);\ \psi_1^* = \rho_1.$$

iii. With high $\rho_1 \in (H_1, 1]$, τ_{UC}^* is always the equilibrium tax plan. The opportunistic type mimics for sure if $\beta_{op} \geqslant 1 - \tau_{UC}^*$ and confiscates for sure if $\beta_{op} \leqslant 1 - \tau_{UC}^*$.

Household production is $\mu_1^* = [(1 - \tau_1^*)\psi_1^*]^\gamma$ in all cases. L_0, L_1, M_1, H_1, A_1 and A_2 are functions of $\beta_{tr}/\beta_{op}, \beta_{op}$ and γ. Their specific forms are detailed in the appendix.[20]

Proof. First of all, $\bar{\tau}_1$ is no longer an optimal choice for the trustworthy type because it becomes an minimizer in this case.[21] The trustworthy type thus chooses between pooling and separating outcomes unless the tax rate hits the feasibility constraint: $[0, 1]$. L_1 is the lowest long-term credibility with which a constrained separating outcome can be induced by a non-negative tax rate. L_1 is thus obtained by solving $\hat{\tau}_1(L_1) = 0$.

Within the region $[0, L_1]$, a constrained separating outcome is not feasible, so the choice is between a constrained pooling outcome and a mixed outcome with zero tax rate. L_0 is the cutoff below which the constrained pooling outcome dominates. This is *Property i.* When $\rho_1 \geqslant L_1$, the choice is between a constrained pooling outcome and a constrained separating one. M_1 is the cutoff above which the constrained pooling outcome dominates.

Finally, when ρ_1 is high enough that τ_{UC}^* can support either a separating or a pooling outcome, the choice is between a constrained and an unconstrained outcome. If both outcomes are pooling or separating, the unconstrained outcome is always dominant whenever it is feasible, as proved in Lemma 1. In these cases, the cutoffs are h_1 and h_1^{-1}, respectively. If one outcome is pooling and the other is separating, however, $\tilde{H}_1$ is the cutoff below which a constrained pooling outcome dominates, whereas $\hat{H}_1$ is the cutoff below which a constrained separating outcome dominates. Therefore, we obtain H_1 as specified in the appendix.

Cutoffs L_0, L_1, M_1, H_1 in the credibility space are functions of β_{op}, γ and β_{tr}/β_{op}, so all three of these factors play important roles in determining the credibility regions. Figure 3 plots L_0, L_1, M_1, $\tilde{H}_1$, $\hat{H}_1$ on the space of $(\rho_1, \beta_{tr}/\beta_{op})$. The left-hand panel depicts the case where $\beta_{op} \leqslant 1 - \tau_{UC}^*$ and the right-hand panel $\beta_{op} \geqslant 1 - \tau_{UC}^*$. From the plots, we can also view L_0, M_1, H_1 as cutoffs in the space of β_{tr}/β_{op} as functions of $\{\rho_1, \beta_{op}, \gamma\}$. There are some noticeable features common to both cases:

[20] Notice that $A_2(\beta_{op}, \gamma) \geqslant (\gamma + 1)/\gamma$ if and only if $\beta_{op} \leqslant 1 - \tau_{UC}^*$, with the equality holding at $\beta_{op} = 1 - \tau_{UC}^*$. Also, when $\beta_{op} \geqslant 1 - \tau_{UC}^*$, $A_2(\beta_{op}, \gamma) \geqslant 1/\beta_{op} \geqslant \beta_{tr}/\beta_{op}$, which is proved in the appendix. This excludes the possibility that a constrained separating region can be next to an unconstrained pooling region.

[21] Proof is in the appendix.

1) $\tilde{H}_1$ is smaller than M_1 for all ρ_1 and is tangential to M_1 at $\rho_1 = h_1^{-1}$, because τ_{UC}^* dominates all the other $\tilde{\tau}_1$ that induce pooling outcomes. Similarly, $\hat{H}_1$ is larger than M_1 and is tangential to M_1 at $\rho_1 = h_1$, because τ_{UC}^* dominates all the other $\hat{\tau}_1$ that induce separating outcomes. 2) It is shown in the appendix that $\lim_{\rho_1 \to 1} \tilde{H}_1 = -\infty$ and $\lim_{\rho_1 \to 1} \hat{H}_1 = +\infty$. Therefore, unconstrained outcomes are always dominant if long-term credibility is high enough. This proves *Property iii.* 3) L_0 and M_1 intersects once at $\rho_1 = L_1$ where $\hat{\tau}_1 = 0$.

Using these features together with the plots in Figure 3, we can obtain the results summarized in *Property ii.* When $\beta_{tr}/\beta_{op} \leqslant M_1\left(\rho_1 = L_1, \beta_{op}, \gamma\right) \equiv A_1\left(\beta_{op}, \gamma\right)$, constrained pooling outcomes are dominant unless $\rho_1 \geqslant H_1$. When $\beta_{tr}/\beta_{op} \geqslant \lim_{\rho_1 \to 1} M_1\left(\rho_1, \beta_{op}, \gamma\right) \equiv A_2\left(\beta_{op}, \gamma\right)$, constrained separating outcomes are dominant whenever they are feasible unless $\rho_1 \geqslant H_1$. When β_{tr}/β_{op} is between A_1 and A_2, however, the MPE is "mixed" with zero planned tax rate for low ρ_1, "constrained separating" for moderate ρ_1; and "constrained pooling" for high ρ_1. ■

Propositions 1 and 2 have in common the fact that:

Corollary 1 *If $\beta_{tr} = \beta_{op}$, the trustworthy type never reduces its tax plan below τ_{UC}^* in equilibrium to invest in reputation.*

Only when the trustworthy type is sufficiently more patient than the opportunistic type does a mixed outcome with $\tau_1^* < \tau_{UC}^*$ or a constrained separating outcome become optimal. In addition, the MPE's property is invariant with β_{tr}/β_{op} with extreme values of long-term credibility ρ_1. Only with moderate ρ_1 does high β_{tr}/β_{op} induce separating outcomes in equilibrium.

In contrast to most cases with $\gamma > 1$, however, the MPE with $\gamma \leqslant 1$ often has a policy function $\tau_1^*(\rho_1)$ that is discontinuous when ρ_1 moves from one credibility region to another. Such high sensitivity of optimal tax plans to credibility levels is essentially a reflection of the high sensitivity of the opportunistic type's strategy to tax changes.

The Effect of Time Horizon – Solutions in a T-period Model

This section extends the previous analysis to a T-period model for two purposes. First, as a robustness check, the extension preserves the central property that β_{tr}/β_{op} is the key determinant in the trustworthy type of government's investment in reputation. Second, the extension allows us to obtain the time series implications of the model.

Denote n as the number of periods left until the terminal stage of the game, so $n = 0, 1, ..., T-1$ corresponds to $t = T, T-1, ..., 1$. When T is large enough, we can always divide the game into two parts: $\gamma > n$ and $\gamma \leqslant n$. Thus, the analysis in Section 3 can be applied accordingly.[22]

When there are more future periods, i.e. n is larger, the opportunistic type's incentive to mimic is strengthened. That is $\psi_n(\rho, \tau) \geqslant \psi_m(\rho, \tau)$ if $n \geqslant m$. The strengthened incentive to mimic makes it increasingly harder for the trustworthy type to signal. On the other hand, if the trustworthy type is more patient than the opportunistic type, the difference between types in terms of their future continuation values becomes more prominent when n is larger. In this case, the trustworthy type puts an increasingly larger weight on future continuation values than the opportunistic type, and is thus more willing to signal its type.

[22] The relative sensitivity to tax changes of the weak type compared to households increases if there are more periods left in the game. That is:

$$\frac{\partial \mu_n(\rho, \tau)}{\partial \tau} > 0 \text{ if and only if } n > \gamma.$$

Which force is dominant again critically depends on the ratio of time discount factors β_{tr}/β_{op}. A low ratio $\beta_{tr}/\beta_{op} < B_1$, makes the trustworthy type yield to the mimicking by the opportunistic type and exert no effort to signal. B_1 is defined as

$$B_1 = \begin{cases} (\gamma+1)/\gamma & \text{if } \gamma \geqslant 1 \\ A_1\left(\beta_{op}, \gamma\right) & \text{if } \gamma \leqslant 1. \end{cases}$$

This is the case that will be presented in Propositions 3 and 4. A high ratio $\beta_{tr}/\beta_{op} > B_2$, on the other hand, motivates the trustworthy type to invest heavily in reputation, knowing that it will later obtain a larger tax base. B_2 is defined as a cutoff in β_{tr}/β_{op} at $\bar{n}$ ($\bar{n} \geqslant \gamma > \bar{n} - 1$ if $\gamma > 1$ and $\bar{n} = 2$ if $\gamma \leqslant 1$), above which the MPE is *not* constrained pooling with moderate long-term credibility $\rho_{\bar{n}}$:

$$B_2 > \begin{cases} (\gamma+1)/\gamma & \text{if } \gamma \geqslant 1 \\ A_2\left(\beta_{op}, \gamma\right) & \text{if } \gamma \leqslant 1 \end{cases}.$$

Proposition 5 will present such a case. These propositions will focus on the properties of time series along the equilibrium path instead of the policy function $\tau_n^*(\rho_n)$ at each horizon, because the latter has an overall pattern similar to its counterpart elaborated in Section 3 but with the exact shape dependent on the details of the model.

Proposition 3 *If $\beta_{tr}/\beta_{op} < B_1$ and $\beta_{op} \geqslant 1 - \tau_{UC}^*$,*

i) The MPE in the early part of the game ($n \geqslant \gamma$) is pooling, starting with unconstrained pooling if T is large enough.

ii) The MPE in the later part of the game ($n < \gamma$) depends on the initial long-term credibility ε. With $\varepsilon \in [l_{k+1}, l_k)_{k=1}^{K\in(\gamma-2,\gamma-1]}$, the MPE is pooling at $n > k$ and is mixed at $n \leqslant k$ with

$$\bar{\tau}_n = \frac{n + 1 - n\beta_{tr}/\beta_{op}}{\gamma + 1 - n\beta_{tr}/\beta_{op}}. \tag{13}$$

With $\varepsilon \geqslant l_1$, the MPE is always pooling. l_k is the boundary of two crediblity regions with MPE being "constrained pooling" and "mixed" respectively at $n = k$.

The planned tax rate in equilibrium thus starts with τ_{UC}^* if T is large enough and increases over time until $n \leqslant k$. Afterwards, it decreases over time until the end of the game.

Proof. The detailed proof is left to the appendix. A feature worth noticing is that when $\gamma > 1$, as in the case of $n = 1$, the support of long-term credibility in each period close to the end of the game can be divided into three regions, each associated with a particular form of MPE: "mixed," "constrained pooling" and "unconstrained pooling," respectively. The regions associated with constrained and unconstrained pooling MPE expand to lower long-term credibility as we move backward in the game, i.e. as more periods are left for the government to care about the future. Thus, at a certain, large enough, horizon n, the region with mixed MPE disappears, as does the region with constrained pooling MPE. ■

Figure 4 plots the equilibrium time series $\{\tau_t^*, \rho_t^*, \psi_t^*, \mu_t^*\}_{t=1}^T$ from a numerical example with parameters $\gamma = 5$, $\beta_{tr} = \beta_{op} \geqslant 1 - \tau_{UC}^*$, $T = 20$ and $\varepsilon = 0.01$. To understand the equilibrium paths, it is useful to recall that the gap between the equilibrium tax plan and τ_{UC}^* captures the credibility concerns that the trustworthy type has when it suffers from imperfect credibility.

In this example, T is large enough so that a tax plan τ_{UC}^* can induce the opportunistic type to mimic for sure. Therefore, the game starts with an unconstrained pooling regime and with the trustworthy type keeping the tax plan at the unconstrained optimal rate τ^*_{UC}. In this regime, the economy behaves as if the trustworthy type has full credibility. As time elapses, fewer future periods remain to reward the opportunistic type for maintaining its reputation. In order to sustain a pooling outcome, the trustworthy type has to reduce the opportunistic type's temptation to confiscate. This is done by raising the tax plan from the unconstrained optimal rate τ^*_{UC}, because not only the gap in actual tax rates between mimicking and confiscating will shrink, but the production base μ will also drop, which will further decrease the gain from confiscation. Such a credibility concern for the trustworthy type becomes severer as we move towards the end of the game, so the equilibrium tax plan departs more and more from τ^*_{UC} and the production base continues to drop.

When it becomes too expensive for the trustworthy type to support a pooling outcome, the mixed regime takes over, with the opportunistic type randomizing between mimicking and confiscating. The randomization makes long-term credibility grow over time whenever the tax plan is implemented since such an event is now informative about the government's type. The growth in long-term credibility in turn mitigates the credibility concern in the optimal policy design. As a result, the equilibrium tax rate declines over time toward the unconstrained optimal rate τ^*_{UC}. Short-term credibility also declines in this case as the opportunistic type's incentive to mimic fades away and its long-term credibility grows slowly. Driven by the declining short-term credibility, the production base μ continues to drop until the end of the game.

When $\beta_{op} < 1 - \tau^*_{UC}$, the equilibrium time series are similar if initial long-term credibility is low, except that an unconstrained pooling outcome is no longer feasible. Instead, an unconstrained separating outcome becomes optimal at each horizon n if long-term credibility is high. The presence of such a credibility region with "unconstrained separating" MPE alters the property of MPE for other levels of long-term credibility with respect to the case where $\beta_{op} \geqslant 1 - \tau^*_{UC}$:

Proposition 4 *If $\beta_{tr}/\beta_{op} < B_1$ and $\beta_{op} < 1 - \tau^*_{UC}$, the equilibrium time series depends on the initial long-term credibility ε:*

i) With low ε, the MPE is constrained pooling at $n \geqslant \gamma$ and is mixed with $\bar{\tau}_n$ defined by (13) at $n < \gamma$. The equilibrium path of the planned tax rate is thus the same as the one in Proposition 3.

ii) With moderate ε, the MPE is mixed first and unconstrained separating later, with the planned tax rate decreasing over time to τ^*_{UC}.

iii) With high ε, the MPE is unconstrained separating throughout the whole game.

Proof. Details of the proof are again in the appendix. However, it is worthwhile to point out that the forms of MPE for credibility regions in this case differ significantly from those in the case where $\beta_{op} \geqslant 1 - \tau^*_{UC}$. More specifically, at each horizon $n \geqslant \gamma$, there are three credibility regions, ranked by ρ from low to high, associated with "constrained pooling," "mixed" and "unconstrained separating" MPE, respectively. At each horizon $n < \gamma$, the region with constrained pooling MPE is replaced by a region with mixed MPE with a particular tax rate $\bar{\tau}_n$ defined by (13). As we move towards the end of the game,.this region with mixed MPE with $\bar{\tau}_n$ expands to higher long-term credibility. But the region with unconstrained separating MPE remains the same throughout the whole game. ■

The common message from Propositions 3 and 4 is that when $\beta_{tr}/\beta_{op} < B_1$, all planned tax rates in equilibrium, τ^*_n, are larger than τ^*_{UC}. Notice also that $B_1 > 1$ at any level of elasticity of supply γ.[23] Therefore, Corollary 1 in Section 3 is robust to any finite-horizon game:

[23] $\min A_1 \left(\beta_{op}, \gamma\right) = \lim_{\rho_1 \to 0} M_1 \left(\rho_1, \beta_{op}, \gamma\right) = \gamma + 1.$

If $\beta_{tr} = \beta_{op}$, the trustworthy type never reduces its tax plan below τ^*_{UC} in equilibrium to invest in reputation.

This result can be understood by the amount that various government types are willing to pay in terms of current tax revenues for their reputation. If the trustworthy type shares a similar time preference with the opportunistic type, the short-run cost necessary to induce confiscation by the opportunistic type does not pay off for the trustworthy type in the long run, either. If the trustworthy type is sufficiently more patient, however, it can endure more current pain to inflate the short-run temptation for the opportunistic type to confiscate, and thus gain more credibility for higher tax revenues in the future. This is the case where investment in reputation by the trustworthy type can be optimal:

Proposition 5 *If $\beta_{tr}/\beta_{op} > B_2$, the equilibrium time series depends on the initial long-term credibility ε:*

i) When ε is not extremely low, the MPE is mixed in the early part of the game and is separating in the later part. The planned tax rate in equilibrium starts with a zero or low rate ($\tau^*_n < \tau^*_{UC}$) and increases over time to τ^*_{UC}.

ii) When ε is extremely low, the MPE in the early part of the game ($n \geqslant \gamma$) is constrained pooling with the planned tax rate increasing over time. After $n < \gamma$, the MPE is mixed first and separating later, with the planned tax rate starting below τ^*_{UC} and increasing over time to τ^*_{UC}.

Proof. Details of the proof are again in the appendix. Recall that this is the case where the increasing incentive to signal dominates the increasing cost with more future periods left in the game, i.e. a larger n. Therefore, when we move backwards to the beginning of the game, the trustworthy type is willing to sacrifice more current tax revenues to accumulate reputation, which implies a lower planned tax rate. However, if long-term credibility is extremely low and the number of future periods in the game is large, the production base will be too small for the opportunistic type to confiscate and give up the future streams of tax revenues. A constrained pooling outcome is thus the equilibrium until the later part of the game, when the opportunistic type is less concerned about maintaining it reputation. Only by then can the trustworthy type find it both feasible and desirable to invest in reputation using a planned tax rate lower than τ^*_{UC}. ■

Using the same parameterization as the previous numerical example, except for decreasing β_{op} so that $\beta_{tr}/\beta_{op} > (\gamma + 1)/\gamma$, I plot the equilibrium time series in Figure 5 as a contrast to Figure 4 where both government types are equally patient.

With a low but not extremely low level of initial credibility, the game starts with a mixed regime where the tax rate is kept at zero by the trustworthy type to invest in its reputation. A zero tax plan generates zero tax revenue when it is enacted. This high loss gives the plan low credibility in the initial periods of the game when the government is likely to be opportunistic, since there is a very high probability that the opportunistic type will deviate from the tough plan. As households have little confidence in being taxed at rate zero, not many of them decide to produce, i.e, the initial production base μ remains low. However, when tough tax plans are implemented consecutively for a number of periods, the evidence that the current government is a trustworthy type quickly accumulates. With more trust in the government's behavior, the production base rises.

After long-term credibility reaches a certain level, it can be maintained with lower investments. This takes place when the trustworthy type starts to raise the planned tax rate to reduce revenue losses. As households are more convinced that the current government is trustworthy, the planned tax rate becomes more credible. Hence, despite higher tax rates, the production base continues to rise and in turn increases tax revenues substantially. Long-term credibility is still growing, but at a lower rate than before.

The mixed regime with positive taxes will last until the trustworthy government decides to distinguish itself absolutely from the opportunistic type. Such an event will occur before the final period if the long-term credibility accumulates to a high point where the revenue cost of implementing such a separation is justifiable. After gaining full credibility, the trustworthy type can tax at the unconstrained optimal rate, with the production base being at its peak as well. The government then achieves the best outcomes after all its earlier effort to overcome its lack of credibility.

As mentioned at the beginning of this section, with a larger n, the trustworthy type is more willing but also finds it harder to signal its type. Neither of the offsetting forces dominates the other when β_{tr}/β_{op} lies between B_1 and B_2 so the properties of the MPE are less clear in this case. The equilibrium time series largely depend on how the equilibrium path of long-term credibility evolves over time. Nonetheless, if $\gamma > 1$, low initial long-term credibility is always associated with constrained pooling outcomes early in the game, whereas high initial long-term credibility is associated with early separating outcomes. The cutoff value of the initial long-term credibility is determined by the policy function at $\bar{n}$ where $\bar{n} \geqslant \gamma > \bar{n} - 1$.

Optimal Imitation in Announcement

In the last two sections, the optimal tax announcements were derived under a restriction on the signaling strategy of the opportunistic government: it always has to imitate the announcement strategy of the trustworthy type. This restriction and its implied optimal announcements will be revisited in this section, in the context of a within-period signalling game where both the signalling strategies and households beliefs are obtained in equilibrium. In a nutshell, the results are: (1) imitation is indeed the equilibrium signaling strategy of the opportunistic government, so that the restriction is not a binding one; and (2) the optimal announcement derived under the restriction coincides with the unique equilibrium in the signaling game, when a refinement by Mailath, Okuna-Fujiwara and Postelwaite [1993] is applied. I thus complete the solution to the finitely-repeated game, determining both within-period signalling and the intertemporal evolution of credibility.

The signaling game

The within-period signaling game is between the government in place (now as the sender) and households (now as the receivers). At the beginning of each period, the households have a prior belief about the government's type ρ, which is given by the updated credibility arising from the last period's government tax action. The current government, with its type unobserved by households, then sends a public message in terms of an announced tax plan τ.[24] Contemplating this message, households update their beliefs about the current government type $\phi(\tau)$ from ρ in a manner that this section studies in detail.[25] The new long-term credibility then serves as the base for the interaction between household production $\mu(\phi, \tau)$ and the government's confiscation probability $\pi(\phi, \tau)$, if the opportunistic type is in place.

[24] I focus here, as above, on the situation in which an opportunistic government, if present, has not been revealed. There is no signalling game if the type has been revealed.

[25] Note that $\phi(\tau)$ will turn out to be uniformative and hence there will be no updating after the message is received. But we are at present allowing for this *potential* updating.

The pure sequential equilibrium

As is standard in the game theory literature, a suitable equilibrium concept for such a sender-receiver game with incomplete information is the "sequential equilibrium" proposed by Kreps and Wilson [1982].[26] Denoting $m(TR)$ and $m(OP)$ as the signalling strategies for the trustworthy and opportunistic types, respectively, a pure sequential equilibrium in this game can be defined as follows:

Definition 2 *The strategies and beliefs* $\{m(TR), m(OP), \mu(\phi,\tau), \phi(\tau)\}$ *form a pure "sequential equilibrium" within any time period if:*[27],[28]

D.3.1) $m(TR)$ *maximizes the trustworthy type's payoff:*

$$m(TR) = \arg\max_{\tau} \tau\mu(\phi,\tau) + \beta_{tr} W'(\phi/\psi).$$

D.3.2) $m(OP)$ *maximizes the opportunistic type's payoff:*

$$m(OP) = \arg\max_{\tau} [1-\pi(\phi,\tau)] \left[\tau\mu(\phi,\tau) + \beta_{op} V'(\phi/\psi)\right] + \pi(\phi,\tau)\mu(\phi,\tau).$$

D.3.3) $\mu(\phi,\tau)$ *maximizes the households' expected payoff:*

$$\mu(\phi,\tau) = G[(1-\tau)\psi],$$

where $\psi = \phi + (1-\phi)[1-\pi(\phi,\tau)]$.

D.3.4) $\phi(\tau)$ *is formed in a Bayesian fashion consistent with the strategies*

$$\phi(\tau) = \Pr(TR|\tau) = \frac{\rho \Pr(\tau|TR)}{\rho \Pr(\tau|TR) + (1-\rho)\Pr(\tau|OP)},$$

where

$$\Pr(\tau|TR) = \begin{cases} 1 & \textit{if } m(TR) = \tau \\ 0 & \textit{otherwise} \end{cases} ; \ \Pr(\tau|OP) = \begin{cases} 1 & \textit{if } m(OP) = \tau \\ 0 & \textit{otherwise.} \end{cases}$$

The equilibrium is always pooling

In Sections 3 and 4, the restriction that the opportunistic government always imitates the announcement of the trustworthy type renders the announced message uninformative. It essentially imposes restrictions on the belief function such that $\phi(\tau) = \rho$ for all τ. However, even when the imitation restriction is removed, any equilibrium of the signaling game must still be a pooled one, in which both types send the same message.

[26] Although sequential equilibrium may be a further refinement of perfect Bayesian equilibrium, the difference between these two concepts is not relevant in the current context. But following the classic papers (Krep and Wilson 1982, Grossman and Perry 1986, Mailath et al 1993) that this work has been built upon, I will continue to use "sequential equilibrium" as the main equilibrium concept.

[27] I leave out the opportunistic type's strategy for its confiscation probability $\pi(\phi,\tau)$ in this definition because: (1) it is not part of the signaling equilibrium as it is neither the sender's signaling strategy nor the receiver's best reaction; (2) it is fully determined by the MPE which I derived in previous sections so long as ϕ replaces ρ; (3) its information is incorporated in the construction of short-term credibility, which enters the payoff functions of both governments and households.

[28] I focus only on pure signaling strategies here. That is, the sender is not allowed to randomize over several messages.

To see this, suppose $m(TR) = \bar{\tau}$ and $m(OP) = \tilde{\tau} \neq \bar{\tau}$. Then, the opportunistic type's identity is perfectly revealed if households observe message $\tilde{\tau}$: that is, $\phi(\tilde{\tau}) = 0$. In turn, the opportunistic type will confiscate for sure as its future credibility will never grow again. Anticipating this confiscation, no household will produce, which implies a zero payoff to the opportunistic type. Given this consequence of sending a message $\tilde{\tau} \neq m(TR)$, the opportunistic type can always improve its gains by deviating to the same message as the trustworthy government. Hence, distinct messages from the trustworthy and opportunistic types can never occur in the equilibrium. In other words, any equilibrium in this signalling game involves the opportunistic government imitating the trustworthy type's message-sending. Therefore, in equilibrium, $m(TR) = m(OP) = \bar{\tau}$ and $\phi(\bar{\tau}) = \rho$.

Multiple pooling equilibria

Thus, removing the previous restriction on the opportunistic type's signaling strategy does not affect any of the equilibrium outcomes. However, one still has to determine the equilibrium message, and the requirement $m(TR) = m(OP)$ leaves the out-of-equilibrium beliefs unspecified. A notorious consequence of such unregulated out-of-equilibrium beliefs is a continuum of sequential equilibria that makes the optimal announcement τ^* derived previously only one of many possible outcomes.

For example, if the households somehow interpret all other messages but $\tau = a$ as indicating the current government is opportunistic, it is indeed optimal for both opportunistic and trustworthy governments to send message a in equilibrium. Therefore, any message $\tau \in [0, 1]$ can be supported in a sequential equilibrium by properly specifying the out-of-equilibrium beliefs.

Disciplining off-equilibrium beliefs

In the face of this indeterminacy problem, many equilibrium refinements have been developed to obtain sharper predictions in various classes of signaling games. The game under study has the feature that it is costless for the opportunistic type to imitate the trustworthy one, which makes many standard refinements problematic.[29] In this class of games, the approach by Mailath, Okuna-Fujiwara and Postelwhaite (1993) turns out to be an appropriate one and further selects a unique equilibrium coincident with the optimal policy prediction derived in Sections 3 and 4.

In their work, Mailath, Okuno-Fujiwara and Postelwaite elaborate Grossman and Perry (1986)'s vision of using Bayesian reasoning to discipline out-of-equilibrium beliefs. Mailath et.al. postulate three key ingredients in constructing belief restrictions. First, any out-of-equilibrium message must be one that is sent in an alternative sequential equilibrium by some set of agents. Second, the incentives that various types of agents have to send an alternative message τ are evaluated by comparison of their benefits in a candidate and an alternative equilibrium. Third, when such comparison induces an out-of-equilibrium probability distribution of sender types computed by Bayes' law, it is then used to generate restrictions on beliefs.

Following Definition 2 in Mailath, Okuna-Fujiwara and Postelwhaite (1993), we can explicitly define this restriction on out-of-equilibrium beliefs, which I label as "strongly coherent out-of-equilibrium belief" to facilitate future reference.[30]

[29] For example, intuitive criterion selects the worst pooling equilibrium in which both types announce a 100 percent tax rate and no household produces. This selection is problematic because it implies a discontinuity in equilibrium outcomes between a game with an infinitesimal chance of a government being opportunistic and a game with no chance of such a type. Another well-known refinement by Grossman and Perry (1986) rules out all equilibria in this game.

[30] This label comes from my parallel work with King and Pasten [2008a] on a related model of costless imitative signalling. There, we also study the out-of-equilibrium beliefs restriction proposed by Grossman and Perry [1986] which

Definition 3 *In an economy with a set of sender types Ω, a candidate equilibrium $\{\hat{m}, \hat{\mu}, \hat{\phi}\}$ and an alternative equilibrium $\{\tilde{m}, \tilde{\mu}, \tilde{\phi}\}$, the message τ gives rise to a "strongly coherent out-of-equilibrium belief" φ about a subset of types if*

D.4.1) $\forall\omega \in \Omega$, $\hat{m}(\omega) \neq \tau$, *and there is a non-empty set* $K = \{\omega \in \Omega | \tilde{m}(\omega) = \tau\}$*;*

D.4.2) $\forall\omega \in K$, $R\left(\left\{\tilde{m}, \tilde{\mu}, \tilde{\phi}\right\}, \omega\right) \geqslant R(\{\hat{m}, \hat{\mu}, \hat{\phi}\}, \omega)$ *and* $\exists\omega \in K$, $R\left(\left\{\tilde{m}, \tilde{\mu}, \tilde{\phi}\right\}, \omega\right) > R(\{\hat{m}, \hat{\mu}, \hat{\phi}\}, \omega)$, *where* $R(\cdot)$ *is the payoff function of the senders;*

D.4.3) $\forall\tilde{\omega} \in K$,

$$\varphi(\tilde{\omega}|\tau) = \frac{\rho(\tilde{\omega}) \Pr(\tau|\tilde{\omega})}{\sum_{\omega\in\Omega} \rho(\omega) \Pr(\tau|\omega)},$$

where $\rho(\tilde{\omega})$ *is the prior belief of type* $\tilde{\omega}$ *and* $\Pr(\tau|\omega)$ *specifies the probability that a sender of type* ω *would issue the message* τ *so that*

$$\Pr(\tau|\omega) = \begin{cases} 1 & \text{if } \omega \in K \text{ and } R\left(\left\{\tilde{m}, \tilde{\mu}, \tilde{\phi}\right\}, \omega\right) > R(\{\hat{m}, \hat{\mu}, \hat{\phi}\}, \omega) \\ 0 & \text{if } \omega \notin K \\ [0,1] & \text{otherwise} \end{cases} \tag{14}$$

The refinement by Mailath et. al. eliminates any candidate equilibrium if it cannot be supported by any strongly coherent out-of-equilibrium belief, when such a belief exists.

Definition 4 (Mailath, Okuno-Fujiwara, and Postelwaite Refinement) *A sequential equilibrium $\{\hat{m}, \hat{\mu}, \hat{\phi}\}$ is "defeated" by an alternative equilibrium $\{\tilde{m}, \tilde{\mu}, \tilde{\phi}\}$ if there exists a type $\tilde{\omega}$ and a message $\tilde{\tau} = \tilde{m}(\tilde{\omega})$ sent by $\tilde{\omega}$ in the alternative equilibrium, such that any strongly coherent belief about message $\tilde{\tau}$ is inconsistent with the supporting belief for the candidate equilibrium:*

$$\varphi(\tilde{\omega}|\tilde{\tau}) \neq \hat{\phi}(\tilde{\omega}|\tilde{\tau}) \text{ for any } \Pr(\tilde{\tau}|\tilde{\omega}) \text{ satisfying (14).}$$

Illustration in a static setting

To see how this equilibrium refinement by Mailath et. al. selects the unique equilibrium τ^*,[31] I first use a static version of the current signaling game, which allows a simple graphical description of the key ideas. I will show in the next subsection that the same ideas apply to the original game.

In the static signaling game, the opportunistic government always confiscates and the payoffs to both types reduce to the momentary tax revenues:

$$R(\tau, \mu, TR) = \tau\mu;\ R(\tau, \mu, OP) = \mu.$$

An indifference curve in (τ, μ) space is thus $\mu = w/\tau$ for the trustworthy type with payoff level w, and $\mu = v$ for the opportunistic type with payoff level v. I draw examples of these indifference curves in Panel A of Figure 6 with dashed lines. In each panel, the horizontal axis is the tax announcement τ and the vertical axis is the participation rate μ. The indifference curve of the trustworthy type $\mu = w/\tau$ is decreasing and strictly convex in τ. Given the announced tax plan τ, higher levels of the trustworthy type's payoff correspond to an upward shift of the indifference curve. By contrast,

we call "weakly coherent belief." Use of these belief definitions allows us to state these two alternative approaches in a common manner, highlighting differences in beliefs.

[31] From now on, I will name a pure pooling equilibrium by the equilibrium message sent in it.

the indifference curve of the opportunistic type is independent of the announced plan τ so it is the horizontal line in the panel. A higher participation rate μ increases the payoff to the opportunistic type, and in turn shifts its indifference curve upwards.

To facilitate the drawing of the graph, I further assume the elasticity of supply $\gamma = 1$ so that

$$\mu(\phi, \tau) = (1 - \tau) \phi(\tau).$$

The participation rate depends on the belief ϕ upon receiving different messages τ and thus fully reflects the out-of-equilibrium beliefs, which are the central focus of this section. I use dotted lines in the Figure to draw such belief-based participation functions.

In addition, the continuum of pooling sequential equilibria can be captured by the solid line $\mu = (1 - \tau) \rho$ in all the panels, where ρ is the prior belief in the current government being trustworthy. Any point a on this solid line can be an equilibrium by properly specifying the out-of-equilibrium beliefs. One example is shown in Panel B, where the households believe that any $\tau \neq a$ is sent by the opportunistic government: $\phi(a) = \rho$ and $\phi(\tau \neq a) = 0$.

The solid line is also the restriction on signaling strategy, which we imposed to derive the optimal announcement τ^* in Sections 3 and 4 as $\phi(\tau) = \rho$ for all τ. Hence, Panel A shows the determination of τ^* in the static context where $\tau^* = \arg\max \tau\mu$ subject to $\mu = (1 - \tau) \rho$. In the discussion below, I will apply the refinement by Mailath et al to rule out any pooling equilibrium other than τ^*, and further establish the "undefeated" feature of τ^* so that it is the unique signaling equilibrium surviving this refinement.

Ruling out high tax equilibria

As shown in Panel C in Figure 6, a candidate equilibrium with a high tax announcement $h > \tau^*$ must be supported by beliefs that make all $\tau \neq h$ less desirable than h for both the trustworthy and opportunistic governments. Reflected in the graph, the belief-implied participation rates must lie below the indifference curves through h for both types. In particular, to avoid the trustworthy government strictly prefering τ^* over h, the supporting belief about τ^* must imply a participation rate lower than $\hat{\mu}$, which is the point on the indifference curve at message τ^*. This requires that: $(1 - \tau^*) \phi(\tau^*) \leqslant \hat{\mu} < (1 - \tau^*) \rho$ and, in turn,

$$\phi(\tau^*) < \rho.$$

However, this belief function is not strongly coherent with the senders' incentives in an alternative equilibrium $\tau^* < h$. For the opportunistic government, it is strictly better off by sending τ^* as its indifference curve through τ^* is above the one through h. For the trustworthy government, deviating to τ^* also strictly increases its payoff because τ^*, by construction, is the unique revenue maximizer when the belief about the equilibrium message is ρ. Therefore, the strongly coherent unconstrained optimal rate receipt of the out-of-equilibrium message τ^* is $\varphi(\tau^*) = \rho$ according to the Bayes' law in D.4.3). In other words, the supporting belief for the candidate equilibrium h such that $\phi(\tau^* < h) < \rho$ is not strongly coherent. Thus, $h > \tau^*$ is *defeated* by the alternative equilibrium τ^*.

Ruling out low tax equilibria

Panel D in Figure 6 shows a candidate equilibrium with a low tax announcement $l < \tau^*$ and its supporting out-of-equilibrium beliefs. Applying the same reasoning as above, to avoid the trustworthy type deviating from l to τ^*, the supporting belief about message τ^* has to satisfy:

$$\phi(\tau^*) < \rho.$$

Again, this belief is not strongly coherent with the senders' incentives in an alternative equilibrium τ^*. The trustworthy government is strictly better off by deviating to τ^* but the opportunistic type is not, since the participation rate in the candidate equilibrium l is higher. So the only strongly coherent belief about τ^*, $\varphi(\tau^*)$, is equal to 1, which is inconsistent with $\phi(\tau^*) < \rho$ – the supporting belief for the equilibrium l. Thus, any candidate equilibrium $l < \tau^*$ is *defeated* by the alternative equilibrium τ^*.

Existence

What remains to be shown is that the candidate equilibrium τ^* survives this refinement – that is, τ^* is undefeated by any alternative equilibrium.

To see this result, first notice that the candidate equilibrium τ^* can be supported by the out-of-equilibrium beliefs

$$\begin{aligned}\phi(\tau) &= 0 \text{ for } \tau < \tau^* \\ \phi(\tau) &= \rho \text{ for } \tau \geqslant \tau^*.\end{aligned}$$

Consider an alternative equilibrium $h > \tau^*$. An opportunistic government does not have an incentive to deviate from τ^* to h because the alternative equilibrium involves a lower participation rate and thus less tax revenues. A trustworthy government will not deviate to h either, since τ^* generates strictly higher revenues in the candidate equilibrium than h does in the alternative equilibrium: $\tau^*(1-\tau^*)\rho > h(1-h)\rho$. If neither government sends message h, no restrictions are placed on the out-of-equilibrium belief about h.

Now consider an alternative equilibrium $l < \tau^*$, which has a higher participation rate so that the opportunistic government will find it desirable to send such a message. However, the trustworthy government will not because τ^* is already the revenue-maximizer given the equilibrium belief ρ. Hence, the strongly coherent belief about the message $l < \tau^*$ is 0, which coincides with the supporting belief for the equilibrium τ^*: $\phi(\tau < \tau^*) = 0$.

Therefore, the equilibrium τ^* is the only survivor of the Mailath, Okuno-Fujiwara and Postelwaite refinement.

Generalization to the dynamic setting

Having displayed the power of this refinement approach in the static setting, I now argue that the same idea works in the original game.[32]

In the dynamic context, the optimal announcement τ^* derived in Sections 3 and 4 is essentially the solution to the following constrained maximization problem after imposing the out-of-equilibrium belief function $\phi(\tau) = \rho$ for all τ:

$$\begin{aligned}\tau^* &= \arg\max_{\tau} \tau\mu(\rho,\tau) + \beta_{tr} W'(\rho/\psi) \qquad (15) \\ \text{subject to } \mu(\rho,\tau) &= G[(1-\tau)\psi].\end{aligned}$$

[32] For readers who are curious about whether this refinement works in a static setting with $\gamma \neq 1$, the answer is yes as long as $\tau^* = \arg\max_\tau \tau(1-\tau)^\gamma \rho^\gamma$ is unique. The reasoning follows exactly that presented in this subsection.

with $\psi(\rho,\tau)$ being determined by $\tau\mu(\rho,\tau)+\beta_{op}V'(\rho/\psi)=\mu(\rho,\tau)$ or its boundary value $\{\rho,1\}$.

There are two complications added by having the governments live more than one period in the signaling game. One is that the payoff functions of governments now include the future continuation values. The other is possible mimicking by the opportunistic type so that its short-term credibility ψ is not necessarily equal to the households' beliefs about the type ϕ. However, neither changes the essence of the equilibrium selection.

Any candidate equilibrium $\tau \neq \tau^*$ is defeated

From the argument in the static setting, the key ingredient in ruling out equilibria other than τ^* is the inconsistency between the supporting out-of-equilibrium belief about message τ^*: $\phi(\tau^*)<\rho$ and the strongly coherent belief about τ^*: $\varphi(\tau^*)\geqslant\rho$. In the dynamic setting, we still need $\phi(\tau^*)<\rho$ as a requirement of the supporting beliefs for any candidate equilibrium $\bar{\tau}\neq\tau^*$, because $\tilde{\phi}(\tau^*)=\rho=\phi(\bar{\tau}\neq\tau^*)$ generates strictly higher payoffs for the trustworthy type by the construction of τ^* so that it will induce the government to deviate from $\bar{\tau}$ to τ^*. Similarly, any belief $\tilde{\phi}(\tau^*)>\rho$ will also induce deviation because the payoff to the trustworthy type always increases with its long-term credibility ϕ.

On the other hand, the strongly coherent belief about τ^*, $\varphi(\tau^*)\geqslant\rho$, is also preserved in the dynamic setting. This stems from the fact that the trustworthy government has a strictly higher payoff in the alternative equilibrium τ^* than in the candidate equilibrium $\bar{\tau}\neq\tau^*$ by the construction of τ^*. Then, at least the trustworthy type will have an incentive to send the out-of-equilibrium message τ^*: $\Pr(\tau^*|TR)=1$. Regardless of the incentive to the opportunistic type, Bayes' law implies $\varphi(\tau^*)\geqslant\rho$.

Therefore, any candidate equilibrium $\bar{\tau}\neq\tau^*$ in the dynamic setting is defeated by the alternative equilibrium τ^* as its supporting belief cannot be strongly coherent with the governments' incentives to send message τ^*, and hence all equilibria $\bar{\tau}\neq\tau^*$ must be discarded from the equilibrium set.

The candidate equilibrium τ^* is undefeated

The undefeated nature of the candidate equilibrium τ^* in the static setting stems from two facts: (1) such an equilibrium can be supported by any belief $\phi(\tau)\leqslant\rho$ for all τ with $\phi(\tau^*)=\rho$; and (2) the strongly coherent beliefs about any message $\tau\neq\tau^*$ are consistent with this class of supporting beliefs.

Both key facts hold in the dynamic setting. For the first, because τ^* is the unique payoff-maximizer for the constrained optimization (15) under the restriction $\phi(\tau)=\rho$ for all τ, the payoff to the trustworthy type for sending τ^* is strictly higher than the payoff for sending $\tau\neq\tau^*$. As the payoff increases with long-term credibility ϕ, sending τ^* is certainly the best choice for the trustworthy type when $\phi(\tau)\leqslant\rho=\phi(\tau^*)$.

For the second fact, by the construction of τ^*, the trustworthy type has no incentive to send any message $\tau\neq\tau^*$ because its payoff in the candidate equilibrium τ^* is strictly higher than in any alternative equilibrium. Hence, for messages $\tilde{\tau}$ which only the opportunistic type wants to deviate to, the strongly coherent belief is $\varphi(\tilde{\tau})=0$, consistent with the supporting beliefs for the equilibrium τ^*. For messages that even the opportunistic type does not want to send, there are no strongly coherent beliefs defined, and hence no restriction imposed on the supporting beliefs.

All in all, the candidate equilibrium τ^* is undefeated by any alternative equilibrium $\tau\neq\tau^*$ in the dynamic setting so that it is the unique signaling equilibrium surviving the refinement by Mailath, Okuna-Fujiwara and Postelwaite [1993].

Conclusions and Remarks

The paper has presented a simple reputation game where a trustworthy government can use policy announcements as an instrument to accommodate public doubts concerning its ability to precommit, or to facilitate public learning of its true identity. The unique Markov perfect equilibrium of the game reveals that it is optimal for a trustworthy government to separate itself from an opportunistic government only when the trustworthy type is sufficiently more patient. If this condition does not hold, it will be too expensive for a trustworthy type to induce a separating outcome through active policymaking and, thus, reputation building will never occure in equilibrium. This partially explains why optimal reputation building by a trustworthy government, which is often seen in practice, has not been adequately addressed in the literature since both types have usually been assumed to be equally patient.

In fact, assuming a trustworthy type to be more patient than an opportunistic type is more natural if we consider government type to be an endogenous choice based on time preference. Suppose that it is costly to obtain a commitment device, then only a more patient government will pay the cost and become trustworthy because, as the model predicts, it will have more to gain from being able to commit.[33]

In sum, this paper provides a transparent framework for thinking about optimal investment in reputation by a trustworthy government, which has been neglected by the existing literature but is often a central issue in actual policymaking. It thus opens the door to new prospects for productive theoretical research in other richer environments. In addition, as this framework involves a consistent Bayesian perspective on choice, signalling and equilibrium selection, some results may ultimately provide a basis for empirical analysis along Bayesian lines.

References

[1] Backus, David A, and John Driffill. (1985a) "Inflation and Reputation." *American Economic Review*, 75:3, 530-38.

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

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

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

[5] Celentani, Macro and Wolfgang Pesendorfer, (1996) "Reputation in dynamic games", *Journal of Eonomic Theory* 70, 109-132

[6] Chari, V. V. and Patrick J. Kehoe, (1990) "Sustainable Plan", *Journal of Political Economy* 98, No.4, 783-802.

[7] Chari, V. V. and Patrick J. Kehoe, (1989) "Sustainable Plan and Debt", working paper, Federal Reserve Bank of Minneapolis.

[33] Note that the comparison is between one scenario where no type buys the commitment device so that the private sector produces nothing in the MPE, and another scenario where governments with certain time preferences buy the commitment device but the time preference is private information.

[8] Chari, V. V., Lawrence J. Christiano and Martin Eichenbaum, (1998) "Expectation Traps and Discretion", *Journal of Economic Theory* 81, 462-492.

[9] Chatterjee, Satyajit, Dean Corbae and Jose-Victor Rios-Rull, (2009) "A Theory of Credit Scoring and Competitive Pricing of Default Risk", mimeo.

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

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

[12] Crawford, Vincent P. and Joel Sobel, (1982) "Strategic Information Transmission", *Econometrica* 50, No.6, 1431-1451.

[13] Cripps, Martin W., George J. Mailath and Larry Samuelson, (2004) "Imperfect Monitoring and Impermanent Reputations", *Econometrica* 72, No.2, 407-432.

[14] Cukierman, Alex, (1993) *Central Bank Strategy, Credibility, and Independence*, MIT press.

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

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

[17] D'Erasmo, Pablo, (2007) "Government Reputation and Debt Repayment in Emerging Economics", mimeo, UTAustin.

[18] Diamond, Douglas W., (1989) "Reputation Acquisition in Debt Markets", *The Journal of Political Economy* 97, No.4, 828-862.

[19] Doraszelski, Ulrich and Mark Satterthwaite, (2007) "Computable Morkov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity", working paper, Harvard Univeristy.

[20] Fischer, Stanley, (1990) "Rules versus Discretion in Monetary Policy", in B. Friedman and F. Hahn, (eds), Handbook of Monetary Economics, North-Holland.

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

[22] Fudenberg, Drew and David K. Levine, (1989) "Reputation and Equilibrium Selection in Games with a Patient Player", *Econometrica* 57, No.4, 759-778.

[23] Fudenberg, Drew and David K. Levine, (1992) "Maintaining a Reputation when Strategies are Imperfectly Observed", *The Review of Economic Studies* 59, No.3, 561-579

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

[25] Horner, Johannes, (2002) "Reputation and Competition", *The American Economic Review* 92, No.3, 644-663

[26] King, Robert G., Yang K. Lu and Ernesto Pasten (2008a) "Coherent beliefs with costless imitative signalling", work in progress.

[27] King, Robert G., Yang K. Lu and Ernesto Pasten (2008b) "Managing Expectations", *Journal of Money, Credit and Banking*, forthcoming.

[28] Klein, Paul, Per Krusell and Jose-Victor Rios-Rull, (2006) "Time-Consistent Public Policy", working paper, Princeton University.

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

[30] Kreps, David M. and Robert Wilson, (1982b) "Reputation and Imperfet Information", *Journal of Economic Theory* 27, 253-279.

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

[32] Liu, Qingmin (2009) "Information Acquisition and Reputation Dynamics", mimeo, UPenn.

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

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

[35] Maskin, Eric and Jean Tirole, (1992) "The Principal-Agent Relationship with an Informed Principal, II: Common Values", *Econometrica* 60, No.1, 1-42

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

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

[38] Myerson, Roger B., (1983) ""Mechanism Design by an Informed Principal", *Econometrica* 51, No.6, 1767-1797.

[39] Ireland, Peter N., (1999) "Expectations, Credibility, and Time-Consistent Monetary Policy", NBER working paper.

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

[41] Rogoff, Kenneth, (1985) "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100, 1169-1190.

[42] Sleet, Christopher and Sevin Yeltekin, (2007) "Recursive monetary policy games with incomplete information", *Journal of Economic Dynamics and Control* 31, 1557-1583.

[43] Sobel, Joel, (1985) "A Theory of Credibility", *The Review of Economic Studies* 52, No.4, 557-573.

[44] Walsh, Carl E., (2003) *Montary Theory and Policy*, MIT press.

Appendix

Proofs for Proposition 2

Parameter definitions in Proposition 2

$$L_0 \text{ solves } \frac{\beta_{tr}}{\beta_{op}}\tau_{UC}^* = \frac{\left[1-\left[\beta_{op}\left(1-\tau_{UC}^*\right)^{\gamma} L_0^{\gamma}\right]^{1/(\gamma+1)}\right]\left[\beta_{op}\left(1-\tau_{UC}^*\right)^{\gamma} L_0^{\gamma}\right]^{(\gamma-1)/(2\gamma+2)}}{1-\left[\beta_{op}\left(1-\tau_{UC}^*\right)^{\gamma} L_0^{\gamma}\right]^{1/2}}$$

$$L_1 = \beta_{op}^{1/\gamma}\left(1-\tau_{UC}^*\right)$$

$$M_1 \text{ solves } \frac{\beta_{tr}}{\beta_{op}}\tau_{UC}^* = \frac{M_1^{\gamma^2/(\gamma+1)} - M_1^{\gamma/(\gamma+1)}}{\left[\beta_{op}\left(1-\tau_{UC}^*\right)^{\gamma}\right]^{1/(\gamma+1)}\left(1-M_1^{\gamma}\right)} + 1$$

$H_1 =$

	$\beta_{tr}/\beta_{op} \leqslant A_2\left(\beta_{op},\gamma\right)$	$\beta_{tr}/\beta_{op} \geqslant A_2\left(\beta_{op},\gamma\right)$
$\beta_{op} \leqslant 1-\tau_{UC}^*$	$\tilde{H}_1$	h_1^{-1}
$\beta_{op} \geqslant 1-\tau_{UC}^*$	h_1	$\hat{H}_1$

$$\text{with } \tilde{H}_1 \text{ solving } \frac{\beta_{tr}}{\beta_{op}}\tau_{UC}^* = \frac{\tilde{H}_1^{\gamma}}{1-\tilde{H}_1^{\gamma}}\left[\left[1-\tilde{\tau}_1(\tilde{H}_1)\right]^{-1} - \left(1+\frac{\tau_{UC}^*}{\beta_{op}}\right)\right]$$

$$\text{and } \hat{H}_1 \text{ solving } \frac{\beta_{tr}}{\beta_{op}}\tau_{UC}^* = \frac{1}{1-\hat{H}_1^{\gamma}}\left[\left(1+\frac{\tau_{UC}^*}{\beta_{op}}\right) - \left[1-\hat{\tau}_1(\hat{H}_1)\right]^{-1}\right]$$

$$A_1\left(\beta_{op},\gamma\right) = \frac{\left[\beta_{op}\left(1-\tau_{UC}^*\right)^{\gamma}\right]^{(\gamma-1)/(\gamma+1)} - 1}{1-\beta_{op}\left(1-\tau_{UC}^*\right)^{\gamma}}(\gamma+1)+\gamma+1$$

$$A_2\left(\beta_{op},\gamma\right) = \frac{1-\gamma}{\left[\beta_{op}\left(1-\tau_{UC}^*\right)^{\gamma}\right]^{1/(\gamma+1)}}+\gamma+1$$

$$\tilde{\tau}_1\left(\rho_1\right), \hat{\tau}_1\left(\rho_1\right), h_1 \text{ are defined in (11) and (12)}$$

$\in \bar{\tau}_1$ is the minimizer when $\gamma \leqslant 1$

The idea of this proof is to show that whenever the second order derivative of the trustworthy type's objective is negative, the first order derivative is also negative. Therefore, $\bar{\tau}_1$ that makes the first order derivative equal to zero can not have negative second order derivative and thus not a maximizer.

Denote U_{tr} as the objective function of the trustworthy type at $t=1$:

$$\begin{aligned} U_{tr}\left(\tau_1;\rho_1\right) &= \tau_1\mu_1\left(\tau_1;\rho_1\right) + \beta_{tr}W_2\left(\rho_1/\psi_1\right) \\ &= \tau_1\mu_1\left(\tau_1;\rho_1\right) + \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\beta_{op}V_2\left(\rho_1/\psi_1\right) \\ &= \tau_1\mu_1\left(\tau_1;\rho_1\right) + \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\left(1-\tau_1\right)\mu_1\left(\tau_1;\rho_1\right). \end{aligned}$$

The last equality stems from the fact that the opportunistic type is indifferent between mimicking and confiscating whenever τ_{UC}^* is not the optimal rate at $t=1$.

The first order derivative of U_{tr} is:

$$\frac{\partial U_{tr}}{\partial \tau_1} = \left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\mu_1 + \left[\left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\tau_1 + \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right]\frac{\partial \mu_1}{\partial \tau_1}.$$

The second order derivative of U_{tr} is:

$$\frac{\partial^2 U_{tr}}{(\partial \tau_1)^2} = 2\left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\frac{\partial \mu_1}{\partial \tau_1} + \left[\left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\tau_1 + \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right]\frac{\partial^2 \mu_1}{(\partial \tau_1)^2}.$$

Now let us write μ_1 and $\partial^2\mu_1/(\partial\tau_1)^2$ as functions of $\partial\mu_1/\partial\tau_1$. To do so, we need the explicit form of μ_1:

$$\mu_1 = (1-\tau_1)^\gamma \psi_1^\gamma = (1-\tau_1)^\gamma \left[\frac{\beta_{op}(1-\tau_2^*)^\gamma \rho_1^\gamma}{(1-\tau_1)^{\gamma+1}}\right]^{1/2} = (1-\tau_1)^{\frac{\gamma-1}{2}}\left[\beta_{op}(1-\tau_2^*)^\gamma \rho_1^\gamma\right]^{1/2},$$

so that

$$\begin{aligned}
\frac{\partial \mu_1}{\partial \tau_1} &= \frac{1-\gamma}{2}\frac{\mu_1}{1-\tau_1} = \frac{1-\gamma}{2}(1-\tau_1)^{\frac{\gamma-3}{2}}\left[\beta_{op}(1-\tau_2^*)^\gamma \rho_1^\gamma\right]^{1/2};\\
\frac{\partial^2 \mu_1}{(\partial \tau_1)^2} &= \frac{3-\gamma}{2}\frac{1}{1-\tau_1}\frac{\partial \mu_1}{\partial \tau_1}.
\end{aligned}$$

Plug them back into the derivatives:

$$\begin{aligned}
\frac{\partial U_{tr}}{\partial \tau_1} &= \left[\left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\frac{2(1-\tau_1)}{1-\gamma} + \left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\tau_1 + \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right]\frac{\partial \mu_1}{\partial \tau_1};\\
\frac{\partial^2 U_{tr}}{(\partial \tau_1)^2} &= \left[2\left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right) + \left[\left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\tau_1 + \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right]\frac{3-\gamma}{2(1-\tau_1)}\right]\frac{\partial \mu_1}{\partial \tau_1}\\
&= \left[2 + \frac{\tau_1}{1-\tau_1}\frac{3-\gamma}{2} + \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\left(\frac{3-\gamma}{2} - 2\right)\right]\frac{\partial \mu_1}{\partial \tau_1}.
\end{aligned}$$

Now, if $\partial^2 U_{tr}/(\partial\tau_1)^2 < 0$, given the fact that $\gamma \leqslant 1$ implies $\partial\mu_1/\partial\tau_1 \geqslant 0$,

$$\begin{aligned}
2 + \frac{\tau_1}{1-\tau_1}\frac{3-\gamma}{2} + \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\left(\frac{3-\gamma}{2} - 2\right) &< 0\\
\frac{3-\gamma}{(1-\tau_1)(\gamma+1)} + 1 &< \frac{\beta_{tr}}{\beta_{op}}\tau_2^*.
\end{aligned}$$

If $\partial U_{tr}/\partial\tau_1 < 0$,

$$\begin{aligned}
\left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\frac{2(1-\tau_1)}{1-\gamma} + \left(1 - \frac{\beta_{tr}}{\beta_{op}}\tau_2^*\right)\tau_1 + \frac{\beta_{tr}}{\beta_{op}}\tau_2^* &< 0\\
\frac{2}{1+\gamma} + \frac{\tau_1}{1-\tau_1}\frac{1-\gamma}{1+\gamma} &< \frac{\beta_{tr}}{\beta_{op}}\tau_2^*.
\end{aligned}$$

Notice that

$$\frac{3-\gamma}{(1-\tau_1)(\gamma+1)} + 1 > \frac{2}{1+\gamma} + \frac{\tau_1}{1-\tau_1}\frac{1-\gamma}{1+\gamma}$$

because

$$\frac{3-\gamma}{\gamma+1} > \frac{1-\gamma}{\gamma+1}.$$

Therefore, if $\partial^2 U_{tr}/(\partial\tau_1)^2 < 0$,

$$\frac{\beta_{tr}}{\beta_{op}}\tau_2^* > \frac{3-\gamma}{(1-\tau_1)(\gamma+1)} + 1 > \frac{2}{1+\gamma} + \frac{\tau_1}{1-\tau_1}\frac{1-\gamma}{1+\gamma}$$

implies $\partial U_{tr}/\partial\tau_1 < 0$.

Relationship between $A_2(\beta_{op},\gamma)$, β_{tr}/β_{op} and $(\gamma+1)/\gamma$

We can rewrite A_2 as

$$\begin{aligned} A_2(\beta_{op},\gamma) &= \frac{1-\gamma}{\left[\beta_{op}/\left(1-\tau_{UC}^*\right)\right]^{1/(\gamma+1)}\left(1-\tau_{UC}^*\right)} + \gamma + 1 \\ &= (\gamma+1)\left[1 + \frac{1-\gamma}{\gamma}\frac{1}{\left[\beta_{op}/\left(1-\tau_{UC}^*\right)\right]^{1/(\gamma+1)}}\right]. \end{aligned}$$

It follows immediately that $A_2(\beta_{op},\gamma) \geqslant (\gamma+1)/\gamma$ if and only if $\beta_{op} \leqslant (1-\tau_{UC}^*)$.

Denote A as $\beta_{op}/(1-\tau_{UC}^*)$. $A \geqslant 1$ when $\beta_{op} \geqslant 1-\tau_{UC}^*$. We can rewrite A_2 and $1/\beta_{op}$ as:

$$\begin{aligned} A_2 &= (\gamma+1)\left[1 + \frac{1-\gamma}{\gamma}A^{-1/(\gamma+1)}\right] \\ 1/\beta_{op} &= A^{-1}(1-\tau_{UC}^*)^{-1} = A^{-1}\frac{\gamma+1}{\gamma}, \end{aligned}$$

and their ratio is

$$\begin{aligned} \frac{A_2}{1/\beta_{op}} &= \left[1 + \frac{1-\gamma}{\gamma}A^{-1/(\gamma+1)}\right]A\gamma \\ &= A\gamma + (1-\gamma)A^{\gamma/(\gamma+1)} \\ &= A^{\gamma/(\gamma+1)} + \gamma\left(A - A^{\gamma/(\gamma+1)}\right). \end{aligned}$$

Since $\gamma/(\gamma+1) < \gamma < 1$ and $A \geqslant 1$, the ratio is then greater than 1. Therefore, $A_2 \geqslant 1/\beta_{op} \geqslant \beta_{tr}/\beta_{op}$.

Limits of $\tilde{H}_1$ and $\hat{H}_1$ when $\rho_1 \to 1$

When we view $\tilde{H}_1$ and $\hat{H}_1$ as the cutoffs in β_{tr}/β_{op} and are functions of ρ_1, β_{op} and γ, we can express $\tilde{H}_1$ and $\hat{H}_1$ as the following:

$$\begin{aligned} \tilde{H}_1 &= \left(\frac{1}{1-\rho_1^\gamma} - 1\right)\left[\rho^{-\gamma/(\gamma+1)}\left[\beta_{op}(1-\tau_{UC}^*)\right]^{-1/(\gamma+1)} - \left(1 + \frac{\tau_{UC}^*}{\beta_{op}}\right)\right]/\tau_{UC}^* \\ \hat{H}_1 &= \frac{1}{1-\rho_1^\gamma}\left[\left(1 + \frac{\tau_{UC}^*}{\beta_{op}}\right) - \rho^{\gamma/(\gamma+1)}\left[\beta_{op}(1-\tau_{UC}^*)\right]^{-1/(\gamma+1)}\right]/\tau_{UC}^* \end{aligned}$$

If we can show that

$$\left[\beta_{op}\left(1-\tau_{UC}^{*}\right)\right]^{-1/(\gamma+1)} \leqslant 1+\frac{\tau_{UC}^{*}}{\beta_{op}}$$

Then immediately $\lim_{\rho_1\to 1}\tilde{H}_1 = -\infty$ and $\lim_{\rho_1\to 1}\hat{H}_1 = +\infty$ since $\left(1-\rho_1^{\gamma}\right)^{-1} \to +\infty$.

To prove the inequality above holds, again denote A as $\beta_{op}/\left(1-\tau_{UC}^{*}\right)$. We can rewrite the inequality equivalently as:

$$\begin{aligned} A^{-1/(\gamma+1)}\left(1-\tau_{UC}^{*}\right)^{-1} &\leqslant 1+\frac{\tau_{UC}^{*}}{1-\tau_{UC}^{*}}\frac{1}{A} \\ A^{\gamma/(\gamma+1)} &\leqslant A\frac{\gamma}{\gamma+1}+\frac{1}{\gamma+1} \end{aligned}$$

The right-hand-side is a linear function of A, and the left-hand-side is a concave function of A since $\gamma/\left(\gamma+1\right) < \gamma < 1$. The two sides equal when $A=1$, and the right-hand-side is tangential to the left-hand-side at $A=1$. Therefore, the inequality must hold.

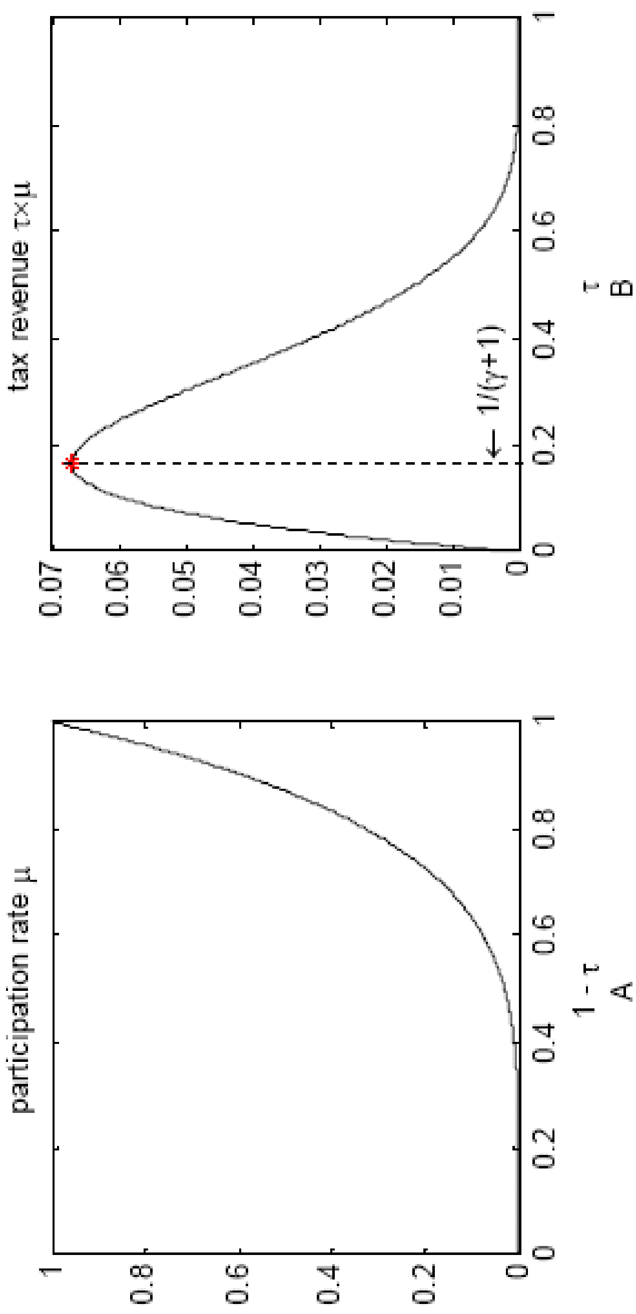

Figure 1: Participation rate and Laffer curve.
The parameter values are: $c=0.4$; $\psi=1$; $\gamma=5$.

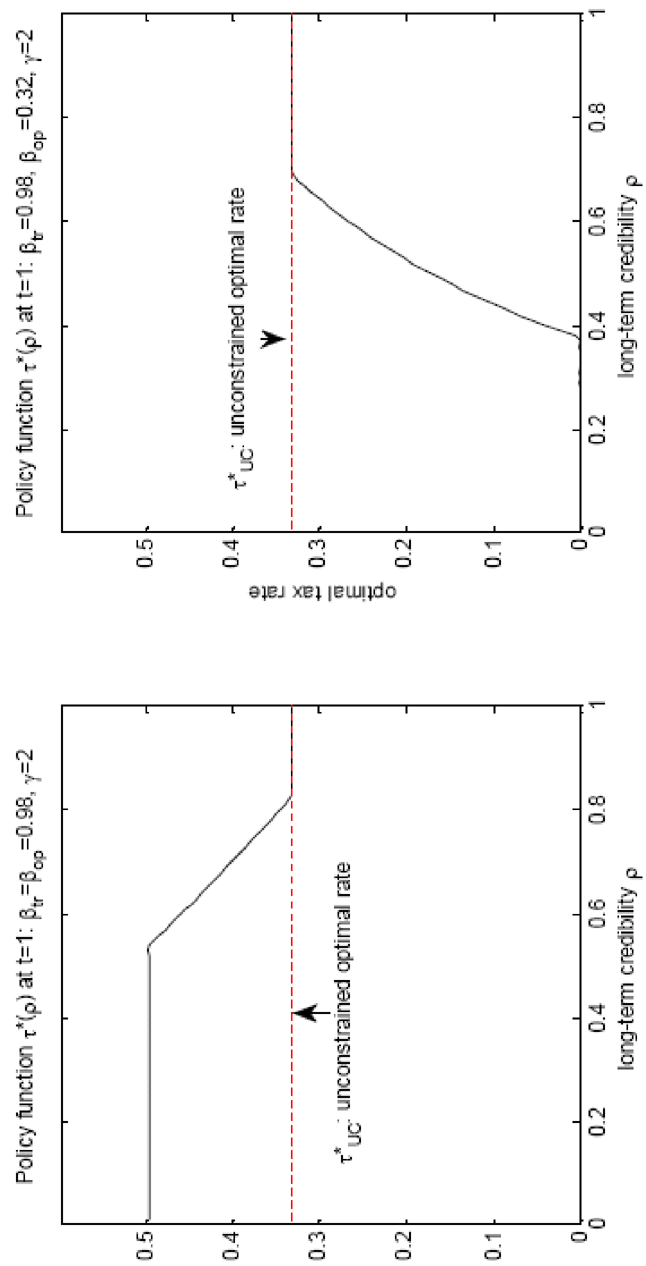

Figure 2: Policy functions at t=1 with γ=2

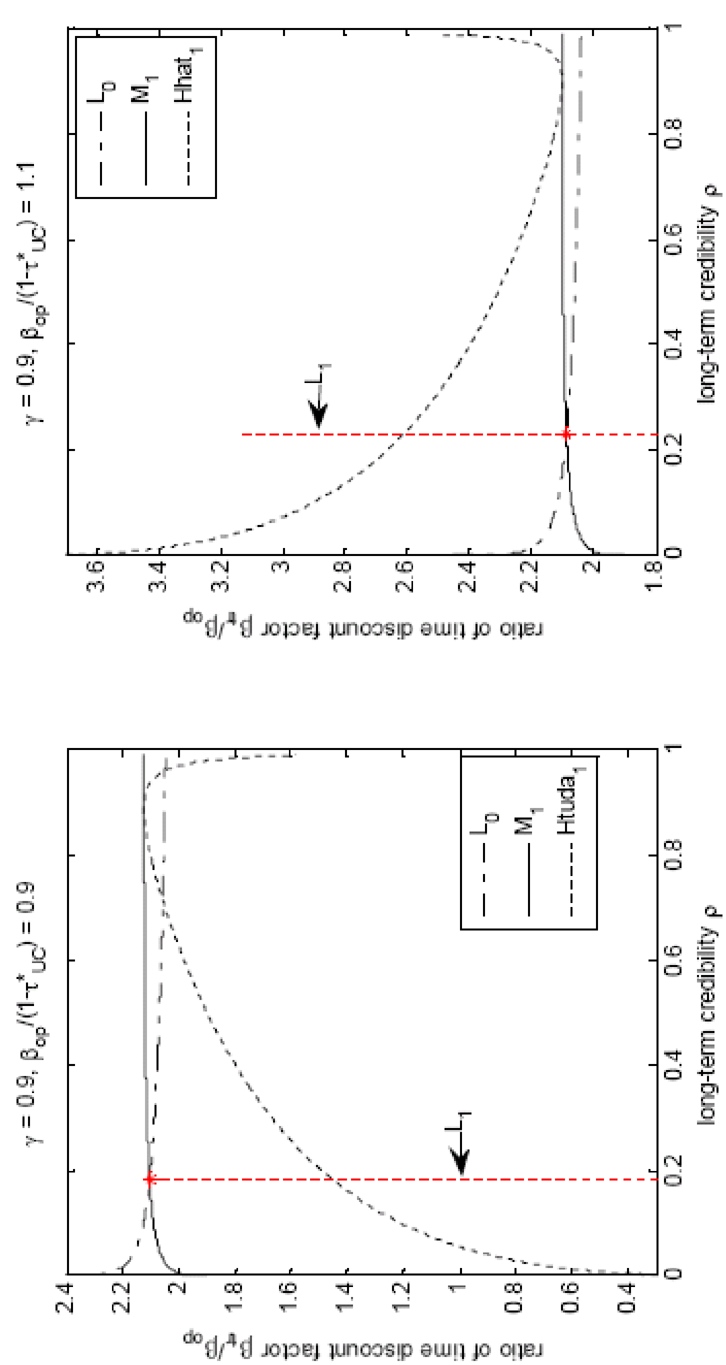

Figure 3: Cutoff values at t=1 with $\gamma<1$

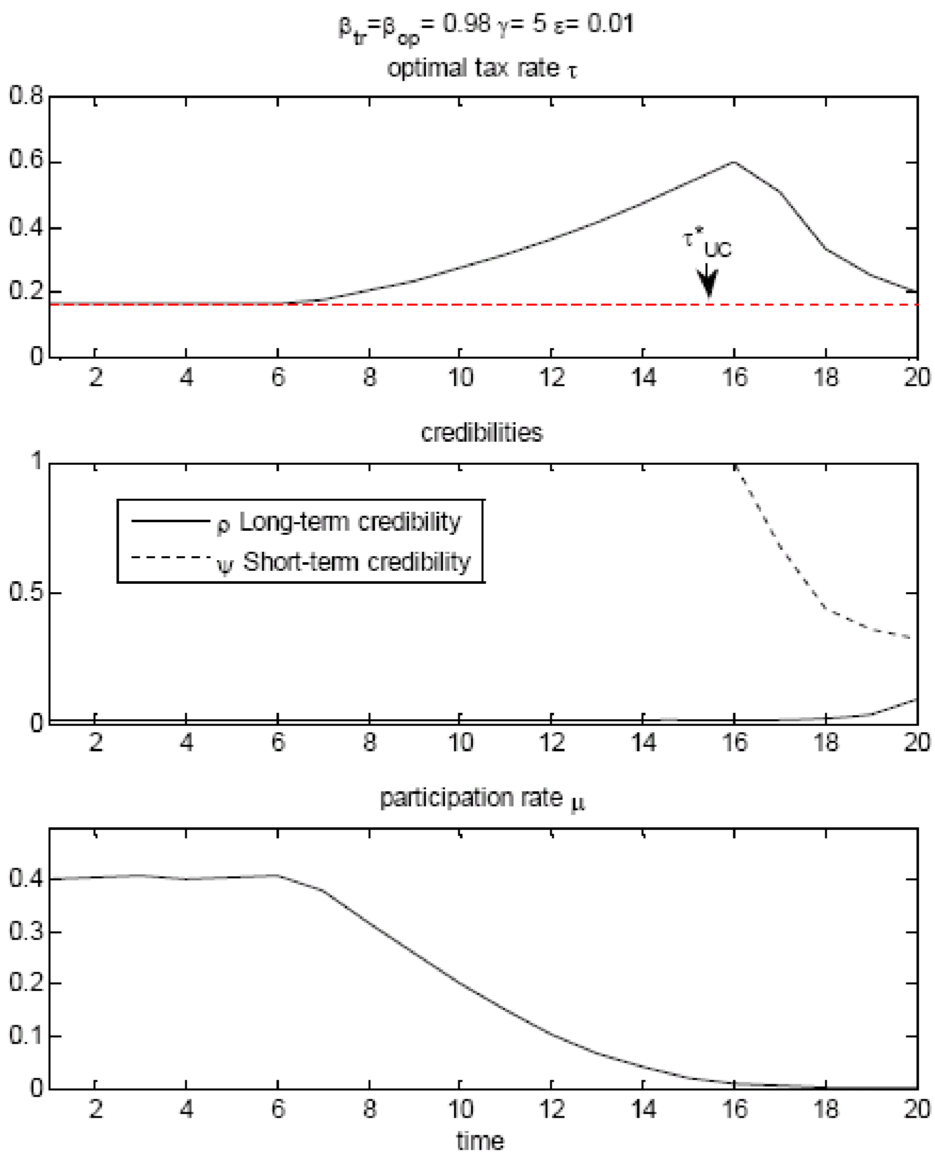

Figure 4: Time series in the case of equally patient government types

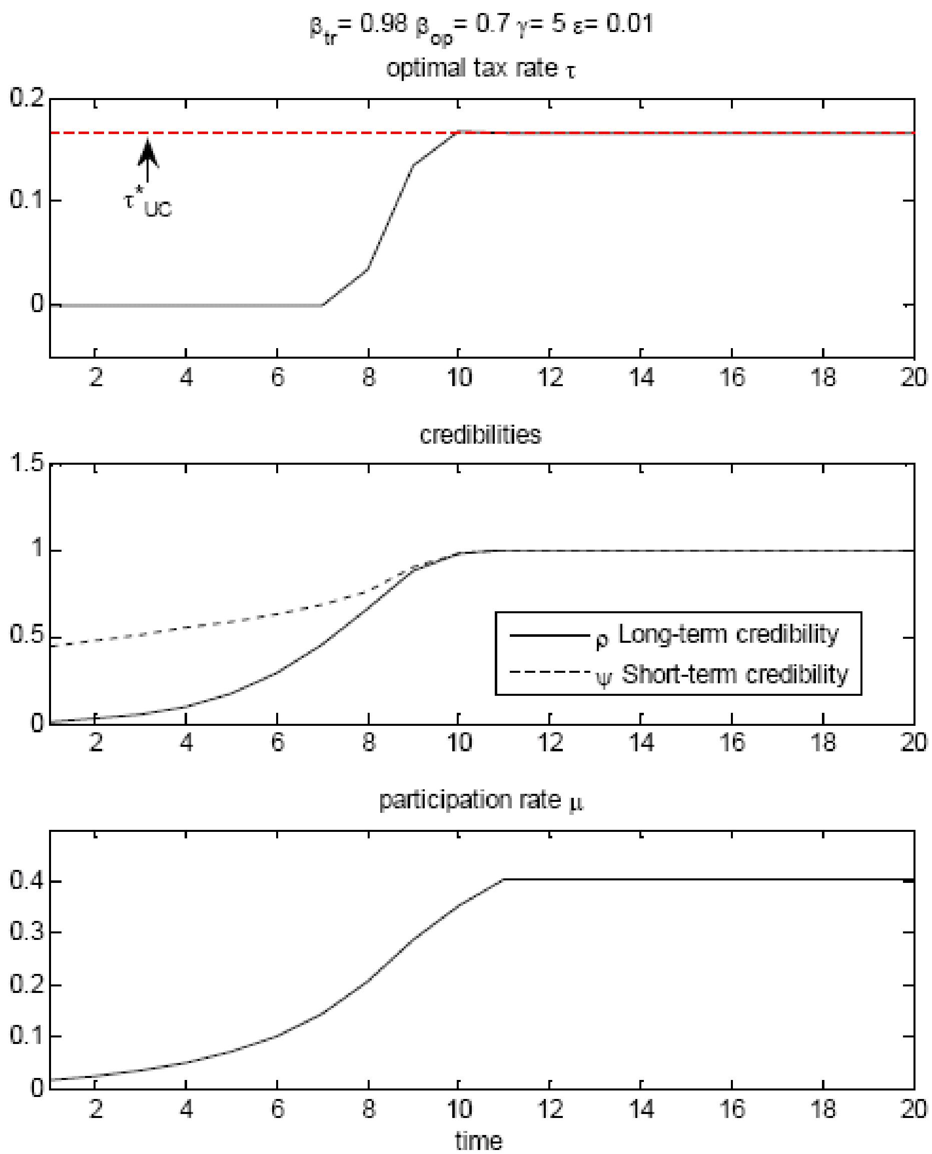

Figure 5: Time series in the case of less patient opportunistic type of government.

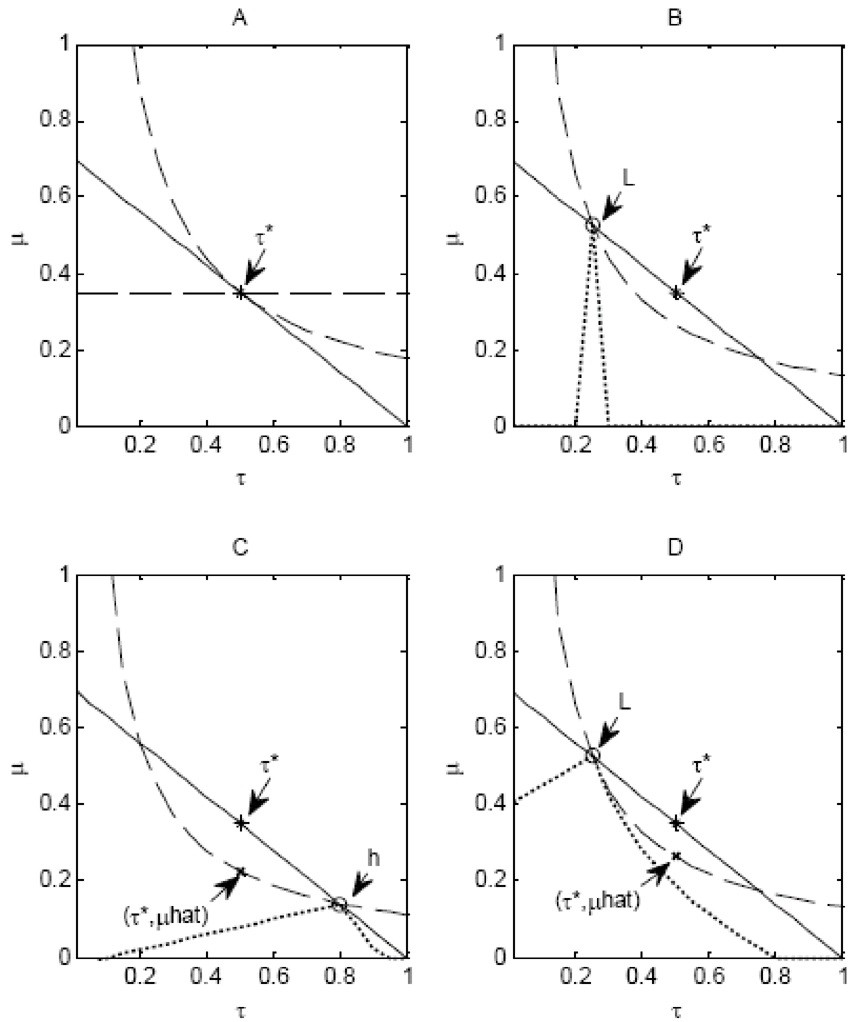

Figure 6: Signaling game in static setting.

European University Institute
European University Press