

Coordination of Expectations and the Informational Role of Policy*

Yang K. Lu[†] and Ernesto Pasten[‡]

June, 2009

Abstract

An *informational role* of policy arises in economies where large fluctuations are triggered by self-fulfilling expectation switches between efficient "optimism" and inefficient "pessimism," a feature that is common in many dynamic economies with coordination failures. Policy affects the information about underlying fundamentals contained in aggregate outcomes, and thus affects the timing of switches and expectations of future switches. We use a problem of optimal taxation on labor income as a laboratory to study this role of policy from a positive and a normative perspective. Our main result is that a stabilization policy is ineffective after an expectation switch. Instead, policy should anticipate switches with small permanent tax cuts to extend "optimism" and severe transitory tax cuts to break "pessimism." These tax cuts should be reverted once a switch is triggered, when policy must focus on its short run objectives.

JEL codes: D8, E6, H2, H3

Keywords: optimal policy, expectation switches, coordination failures, equilibrium selection.

*We specially thank Bob King for his constant support, as well as insightful comments by Elias Albagli, Marios Angeletos, Christophe Chamley, Micael Castanheira, Huberto Ennis, Simon Gilchrist, Francois Gourio, Peter Ireland, Denny Lie, Bart Lipman, Nicola Pavoni, Bernard Salanié, Dan Schulte, Zheng Song, Jean Tirole, Adrien Verdelhan, and seminar participants at Autònoma Barcelona, BU, European Central Bank, Toulouse, 2007 LACEA Meeting (UAndes), the 2008 Far Eastern (SMU), the 2008 European (Bocconi) and the 2009 North American (BU) Meetings of the Econometric Society, ENTER Jamboree (UCL), the 2009 Atelier CIRPEE-DEEP-TSE, Central Bank of Chile, UChile-Econ, UChile-CEA and PUC-Chile. Errors and omissions are exclusively ours.

[†]yanglu@bu.edu. Boston University. 270 Bay State Road, Boston, MA 02215. USA.

[‡]epasten@cict.fr. Toulouse School of Economics and Central Bank of Chile. Manufacture des Tabacs, 21 Allée de Brienne, 31000 Toulouse, France.

1 Introduction

Tracing back to Keynes, self-fulfilling expectation switches between "optimism" and "pessimism" have been considered an important component of booms and slumps in a variety of contexts. Some examples are large fluctuations in output, employment and investment, international capital flows, and certainly financial crises. Pessimism is believed to be inefficient because, if expectations were coordinated in optimism, slumps would be milder and recoveries faster. This paper explores a novel, *informational role* of policy in this class of environments.

We rationalize expectation switches as the result of strategic complementarity among agents that delivers self-fulfilling "optimistic" and "pessimistic" equilibria, heterogeneous information about underlying fundamentals that selects one of these equilibria at each point in time, and shocks that occasionally reveal a good deal of information about these fundamentals. This information triggers synchronized revisions of expectations that provoke switches in the equilibrium selected. Building on this framework, we show that a government with no control and no information advantage about fundamentals still has control on "expectation switches." This result is important, since in most contexts with self-fulfilling dynamics, there is no reason to believe that the government has special knowledge or control of fundamentals of the economy.¹

The informational role of policy arises because the distortionary policy not only affects agents' incentives for their own decisions, but it is also a public signal that affects expectations about others' decisions. These effects imply that policy has control on the total size of the accumulated shocks that is necessary to trigger expectation switches. In turn, policy affects how much agents can learn about fundamentals each period from observing whether or not there is a switch. Such learning effect endows the policy with partial control on future expectation switches. This paper contributes to the understanding of this informational role from both a positive and a normative perspective. In particular, we focus on (1) the channels in which policy affects welfare; (2) the policy trade-offs involved; (3) the distinctive features of the optimal policy with respect to standard policy prescriptions; and (4) the welfare implications of following these standard policy prescriptions.

Specifically, we approach these issues by studying optimal taxation in a dynamic economy without capital or public debt, where labor participation is driven by expectation switches. To build such economy, we apply a mechanism proposed by Chamley [10] to produce regime-switching dynamics and connect it with the Global Games literature. Agents may work or not work, with a reward depending on labor participation (strategic complementarity), and a heterogeneous fixed cost. A single parameter governs the cost distribution, which is called the "fundamental" and is interpreted as the inverse of aggregate productivity. When productivity is extremely low or high, there is a unique equilibrium with respectively low or high labor participations. But when productivity is moderate, both low and high participations may be sustained as self-fulfilling equi-

¹For instance, it is difficult to argue that, in the context of a financial crisis, the government could have better information than market participants about their own liquidity needs or the quality of their assets. Similarly, it is unlikely that the government could have control on technology shocks in a business cycles context.

libria. However, if productivity is unobservable, there is a unique equilibrium characterized by a "switching threshold," such that the equilibrium with low (high) participation takes place when the fundamental is above (below) this threshold. The fundamental is subject to shocks and follows a persistent process, so the level of labor participation in one period provides public information about the state of the fundamental in the next period. Thus, for instance, if participation in one period is low, agents infer that the fundamental is above the switching threshold (productivity is low). Agents then update their expectations accordingly, so in equilibrium the switching threshold is low in the next period. Hence, participation will be low for most realizations of the fundamental so that it is likely to confirm agents' expectations. However, when shocks accumulate such that the fundamental is below this threshold (productivity is high), participation does not coincide with what agents expect. In fact, agents can perfectly infer the fundamental from the level of participation observed. Agents then update their expectations, which now implies a high equilibrium switching threshold. Hence, participation now will be high for most realizations of the fundamental. As a result, aggregate dynamics are characterized by "pessimistic" and "optimistic" regimes with low and high labor participations, with transitions triggered by expectation switches.

We conduct our policy analysis by introducing a benevolent government that taxes labor income to finance a public good. The government has no control on the fundamental, but, similar to agents, it observes past labor participations. The tax rate is announced period by period before agents' decisions. Taxes affect the provision of public goods by the standard Laffer curve. Taxes also affect both the reward of working and agents' expectations about others' participation. Thus, higher taxes not only decrease labor participation in either regime, they also decrease the switching threshold. This latter effect is key for our results: The switching threshold determines the switching probability, which information revealed during a switch and, if there is no switch, how expectations about the fundamental are updated. The pessimistic regime is a dynamic version of a coordination failure since it delivers low reward of working, low labor participation and low public good provision for the same fundamental that can support an optimistic regime.² As a result, the government has incentives to implement a "stabilization policy" that shortens episodes of pessimism and extends episodes of optimism.

Therefore, each of these effects by taxes has an impact on welfare, which makes the government's problem tricky. Recall that the fundamental is above a low switching threshold during pessimism and below a high threshold during optimism, and that higher taxes decrease the threshold. Therefore, a tax cut during a pessimistic regime increases switching probability. However, if the switch is not triggered, it reveals that productivity is low (the fundamental is above a higher switching threshold than if taxes were not cut). This information depresses future expectations, which is a double edge of expansionary policies during pessimism. Conversely, a tax cut in an optimistic regime decreases switching probability. If a switch is triggered, however, agents learn that produc-

²For simplicity, we focus on a case when optimism is efficient. In the case when optimism is inefficient, the channels through which policy affect expectations and the main features of optimal policy are similar to the ones obtained in our analysis.

tivity is very low (the fundamental is above a higher switching threshold than if taxes were not cut). As a result, decreasing taxes extends optimism, but if it fails, it also depresses future expectations.

Moreover, the government could conduct "experiments" by raising taxes, so as to improve agents' and its own knowledge about fundamentals. However, tax experimentation increases the risk of falling into pessimism, so the incentives for such policy decrease with the size of the welfare loss of pessimism.

To study optimal policy, we first impose the assumption that equilibrium uniqueness is preserved under any policy scheme. This allows us to derive a number of results.

Right after an expectation switch, there is a low probability to a new switch in the near term. This is because agents have precise information about the fundamental to be far from the new switching threshold. Hence, policy has no leverage to pull the economy out of pessimism and there is no need for policy to ensure optimism. Thus, optimal policy at the onset of regimes should concentrate on the optimal provision of the public good in the current period. This result stresses the high cost and low return of a standard stabilization policy reacting after large fluctuations.

As the effect of policy on agents' learning is asymmetric across regimes, the optimal policy is also asymmetric. In a pessimistic regime, the gain of a tax cut in raising switching probability *increases* with the size of such a tax cut. This is because, after observing low participation in the period before, agents assign high probability that the fundamental is above its switching threshold. Hence, once the government attempts to break pessimism, it should do so by implementing a severe tax cut (or a "big push")³. However, because expectations get further depressed if this policy fails, there is little incentive to keep the push for two consecutive periods. As a result, policy follows an oscillatory path with occasional big pushes between periods when policy focuses on the short-run public good provision. This result stands against the commonly suggested prescription that policy reactions should be smooth over time.

In contrast, the gain of a tax cut in an optimistic regime in reducing switching probability *decreases* with the size of such a tax cut. This is because, after observing high participation in the period before, agents assign high probability that the fundamental is below the switching threshold. Since tax cuts also involve social costs, a tax cut to delay a switch should be small. If this policy is successful, switching probability continues to increase in the future, which provides incentives for keeping the tax cut. As a result, small but permanent tax cuts are sequentially implemented to extend the optimistic regime, until the government does not want to give up more public goods. Once the switch to pessimism is finally triggered, taxes are raised to attain the optimal provision of the public good, starting a new cycle.

If we relax our assumption that equilibrium uniqueness is preserved for any tax scheme, a successful big push that triggers a switch to optimism may break uniqueness. This is because the level of productivity revealed during the expectation switch is not high enough to unambiguously

³We call this policy a "big push," paraphrasing Murphy, Shleifer and Vishny [23]. These authors propose a "big push" policy to break static a coordination failure in the context of industrialization.

support high participation once taxes are raised in the next period. We propose a fix for this source of multiplicity. A tax profile contingent on which equilibrium takes place serves as an insurance that recovers the dominance of the high participation equilibrium.⁴

This paper is related to the study of equilibrium multiplicity and coordination failures in many economic contexts.⁵ It is also related to Global Games as an equilibrium selection device, which has been mainly applied in static environments.⁶ The selection power of Global Games, however, is fragile when the government controls fundamentals (Angeletos, Hellwig and Pavan [1]) and when asset prices serve as public signals (Angeletos and Werning [4]). But even with equilibrium multiplicity, fundamentals persistence is a force strong enough to generate episodes in which one equilibrium is more likely to take place (Angeletos, Hellwig and Pavan [2]). Dynamics in our economy have the same flavor, but we focus on a case with equilibrium uniqueness where the government has no control and no information superiority about fundamentals, and the only source of public information is the history of aggregate outcomes. Thus, we isolate the effects of policy from this source of information such that this is the only force that can recover equilibrium multiplicity.

Our paper is also related to the empirical work of Hamilton [16] who showed the good fit of regime-switching models in the U.S. post-war data for output.⁷ In a theoretical level, Azariadis [5], Woodford [26] and Howitt and McAfee [19] noticed the potential of coordination failures to conciliate "animal spirits" and rational expectations, but only Chamley [10] and Frankel and Pauzner [15] obtained endogenous mechanisms.⁸ Finally, our paper is also linked to model uncertainty, such as Svensson and Williams [25], who see policy as an experimentation tool in economies following an exogenous regime-switching process. However, we show that once regimes and switches are equilibrium phenomena, such policy motive is undermined.

The paper is organized as follows. Section 2 uses a static Global Games example to preview the main ideas in the paper. Section 3 displays the full dynamic economy. Section 4 introduces the government. Section 5 studies the informational role of policy and Section 6 concludes. An Appendix contains some long proofs skipped in the main text, a numerical example, and figures.

2 The core idea

This section uses a static Global Games example to preview the mechanics and the channels in which the *informational role* of policy operates in an economy with expectation switches.

⁴This contingent design has the potential of fully eliminating inefficiency, but its implementation is unfeasible except on the onset of optimism.

⁵Some examples are Diamond [13], Cooper and John [11], Kiyotaki and Moore [18], and Obsfeld [24].

⁶Proposed by Carlsson and Van Damme [8], it has first been applied to macroeconomics by Morris and Shin [22]. Most dynamic applications assume that the fundamental is revealed at the end of each period.

⁷His results have been confirmed by more recent studies, such as Hamilton and Chauvet [17].

⁸We choose the Chamley's approach instead of Frankel and Pauzner [15] because the later assumes that opportunities to change agents' decisions arrive at a constant rate, which we view as a less flexible assumption.

2.1 A static example

Angeletos and Werning [4] have provided one version of a static example. For such an example, assume that a continuum of agents with total measure one have a binomial decision $a_i \in \{0, 1\}$ and payoffs $U(a_i, A, \theta)$:

	$A \geq \theta$	$A < \theta$
$a_i = 1$	$1 - c$	$-c$
$a_i = 0$	0	0

This game has been used to model speculative attacks against an exchange rate peg, such that $a_i = 1$ is the decision to attack, A is the size of an attack, and θ is the strength of the central bank to resist the attack. We take a different interpretation for a better match with the dynamic economy to be displayed in Section 3, such that $a_i = 1$ is the decision of working, A is labor participation and θ is an exogenous level of participation above which the reward of working is positive. Following Cooper and John [11], a coordination game is characterized by strong strategic complementarity, such that $U_{a,A} > -U_{aa}$ in a range of A . This condition is satisfied here by the discontinuity of the reward of working when $A = \theta$. Hence, if θ is observable, this game has multiple equilibria, $A_L^* = 0$ and $A_H^* = 1$ for $\theta \in [0, 1]$ and $c \in [0, 1]$, where A_H^* is Pareto superior with respect to A_L^* . Thus, the possibility that expectations could be coordinated in A_L^* represents a source of inefficiency without fundamental reasons, known as a "coordination failure."

Global Games modifies the game described above to introduce heterogeneous information among agents about a "fundamental." As shown by Carlsson and van Damme [8] and Morris and Shin [21], the equilibrium under this assumption is unique. Let θ be the fundamental such that

$$\theta \sim N\left(z, \frac{1}{\alpha_z}\right),$$

where z is observable and is interpreted as a source of public information with precision α_z .

Heterogeneity of information is introduced by assuming that each agent receives a noisy signal x_i about θ with precision α_x ,

$$x_i \sim N\left(\theta, \frac{1}{\alpha_x}\right).$$

As in any economy with heterogeneity, the equilibrium is characterized by a cutoff strategy, which is denoted as $x^*(z)$ and depends on public information z . Hence, an agent works if and only if $x_i \geq x^*(z)$. Given x^* , the equilibrium labor participation $A(x^*, \theta)$ is

$$A(x^*, \theta) = \int_{x^*}^{+\infty} d\Phi(x_i | \theta),$$

which depends on θ because the distribution of signals depends on θ . The cutoff $x^*(z)$ is implicitly

defined by the indifference condition of the marginal worker who observes $\{x^*, z\}$:

$$E_\theta [U(1, A(x^*, \theta), \theta) | x^*, z] = E_\theta [U(0, A(x^*, \theta), \theta) | x^*, z]. \quad (1)$$

In this example, a worker receives $1 - c$ if $A(x^*, \theta) \geq \theta$ or, in other words, if $\theta \leq \theta^*(x^*)$, with $\theta^*(x^*)$ defined as $A(x^*, \theta^*) = \theta^*$. Hence, the indifference condition (1) takes the form

$$\Pr[\theta \leq \theta^*(x^*) | z, x^*] = c. \quad (2)$$

This is a fixed point problem which has a unique solution $x^*(z)$ if $\sqrt{2\pi\alpha_x} \geq \alpha_z$. Intuitively, if private signals are more precise than public signals, heterogeneity of information introduces enough heterogeneity of expectations that prevents agents' coordination in an arbitrary equilibrium. As a result, in equilibrium, there are low incentives to work if $\theta > \theta^*(x^*(z))$ and high incentives to work if $\theta \leq \theta^*(x^*(z))$.⁹

If payoffs are continuous in A instead of discontinuous as in this example, the solution of (1) requires the computation of the distribution of $A(x^*, \theta)$ for an agent who observes $\{x^*, z\}$.

2.2 The informational role of policy

Next, our model introduces a government to this game. Such a player has no control and no information advantage about θ but does have control over policy $\{\tau, g\}$ that modifies the payoffs to $U(a_i, A, \theta, \tau, g)$:

	$A \geq \theta$	$A < \theta$
$a_i = 1$	$(1 - \tau) - c + g$	$-c + g$
$a_i = 0$	g	g

with τ representing a distortionary policy and g representing a lump-sum policy. If θ is observable, this game exhibits equilibrium multiplicity for $\theta \in [0, 1]$ and $c \in [0, 1 - \tau]$. Thus, unlike g , τ affects the range of parameters with equilibrium multiplicity.

After introducing heterogeneous information, the equilibrium condition (1) becomes

$$E_\theta [U(1, A^*, \theta, \tau, g) | z, x^*, \tau, g] = E_\theta [U(0, A^*, \theta, \tau, g) | z, x^*, \tau, g],$$

where policy $\{\tau, g\}$ enters in this expression because it affects payoffs and it is a source of public information. In our example this indifference condition becomes

$$\Pr[\theta \leq \theta^*(x^*) | z, x^*] (1 - \tau) = c, \quad (3)$$

⁹This argument survives the limit case when private signals become infinitely precise, $\alpha_x \rightarrow \infty$. In that case $x^*(z) \rightarrow \bar{x}^*$ and $\theta^* = \bar{x}^*$. Thus $A^* = 0$ if $\theta > \bar{x}^*$ and $A^* = 1$ if $\theta \leq \bar{x}^*$.

where $\theta^*(x^*)$ is still defined by $A(x^*, \theta^*) = \theta^*$, unaffected by policy. But the equilibrium cutoff now depends on τ , $x^*(z, \tau)$. Notice that only distortionary policy τ enters in this equilibrium condition because lump-sum policy g by definition does not affect the decision of working.

As the government has no more information than agents about the fundamental θ , shifts of τ do not convey information about θ . However, a higher τ decreases incentives to work, which increases the cutoff x^* and decreases labor participation $A(x^*, \theta)$. A lower labor participation in turn decreases the threshold $\theta^*(x^*)$ that increases x^* even further. Intuitively, there is an amplification effect because a higher τ decreases expectations about what other agents do, which decreases even further the incentives to work because of the strategic complementarity. The overall impact of τ on $x^*(z, \tau)$ affects ex-ante welfare through four channels, which we call the "informational role":

(a) *Momentary welfare.* Since $x^*(z, \tau)$ affects participation $A(x^*(z, \tau), \theta)$, policy τ affects the cross-sectional sum of agents' payoffs regardless of whether the reward of working is $1 - c$ or $-c$.

(b) *Equilibrium selection.* Since $A(x^*(z, \tau), \theta)$ determines the threshold $\theta^*(x^*(z, \tau))$, policy τ affects the likelihood that the equilibrium reward of working is $1 - c$ or $-c$.¹⁰

(c) *Learning.* If agents observe that the payoff of working is $1 - c$, they learn that $\theta \leq \theta^*(x^*(z, \tau))$. And if the reward of working is $-c$, they learn that $\theta > \theta^*(x^*(z, \tau))$. Hence, policy τ affects which information about the fundamental is revealed by realized outcomes. If this game is dynamic and if the fundamental θ_t is persistent, as in our dynamic economy below, the information revealed in previous periods is part of the public information z_t . Hence, current policy τ_t affects future equilibrium cutoffs, and thus affects future momentary welfare and future equilibrium selections.

(d) *Equilibrium multiplicity.* In a dynamic version of this game, the cutoff $x^*(z_t, \tau_t)$ is directly affected by current policy and indirectly affected by past policy through z_t . Therefore, it is possible that some sequences of policy could recover equilibrium multiplicity, which may be or may not be desirable under different circumstances.

One could wonder why there is policy in this game in the first place. Our answer is that, in reality, agents' decisions with the potential of generating self-fulfilling dynamics do interact with policies, such as taxes, monetary policy, or regulatory requirements. These policies typically have objectives unrelated to the management of expectation switches. However, we show that, as long as policy is distortionary, these forms of policy have an informational role. We study this role both from a positive perspective – how policy shifts trigger expectation switches – and from a normative perspective – how to manage the conventional objectives of policy with this informational role. Thus, to study this role, and in particular the dynamic channels in (c) and (d), we need a dynamic environment in which policy is the only force that can recover equilibrium multiplicity. This is the task we turn to in the next section.

¹⁰Ennis and Keister [14] focus on this channel, but they do not use global games to select an equilibrium.

3 A dynamic laboratory economy

This section displays a dynamic economy that serves as a laboratory for our analysis. We model it as a labor economy in which we apply the argument proposed by Chamley [10] to generate endogenous regime-switching dynamics; in this case, on labor participation. Despite the apparent differences between this argument and global games, they share the same intrinsic logic: Chamley's argument also uses heterogeneity in information to produce equilibrium selection when two equilibria are possible. Yet Chamley's argument is dynamic, in which the fundamental is persistent and agents cannot observe the fundamental at the end of each period.¹¹ Hence, past equilibrium realizations become a public signal that affects current equilibrium selection. This effect is what generates regimes with low and high labor participations and switches between them. These switches are the result of synchronized revisions in expectations, so called "expectation switches", which are triggered when shocks occasionally reveal precise information about the fundamental.

Chamley's argument also has a number of other attractive features with respect to global games. First, agents have heterogeneous costs to work, such that equilibrium multiplicity is a combined result of strategic complementarity and a bell-shaped distribution of this heterogeneity. This combined effect alleviates the common critique that coordination games require implausible degrees of strategic complementarity to motivate equilibrium multiplicity. Second, the fundamental is a parameter that governs the distribution of this heterogeneity, which can be interpreted as aggregate conditions, such as productivity. Third, if agents do not observe the fundamental, their own characteristics play the role of noisy signals. We see this motivation for heterogeneity of information as more natural than unmodelled private signals. And finally, some simplifying assumptions deliver unique equilibrium dynamics, allowing us to keep the transparency of our policy analysis in sections 4 and 5. In particular, agents are short-lived, and specific functional forms are assumed for agents' heterogeneity and the stochastic process of the fundamental. Hence, policy shifting is the only force that can recover equilibrium multiplicity.

3.1 Set-up

Players. There is an infinite sequence of generations of agents with measure one indexed by $i \in [0, 1]$. Each agent lives one period, owns one indivisible unit of labor, and takes a binary decision: to work or not to work.

Payoffs. If an agent i living at t decides to work, her payoffs are the difference between her participation reward and a fixed cost. If she decides not to work, she receives utility from leisure

¹¹Most dynamic applications of global games assume that the fundamental is observable at the end of each period.

or home production, which is normalized to zero. Hence, the utility function takes the form

$$U(a_{it}, A_t, c_{it}) = \begin{cases} m(A_t) - c_{it} & \text{if } a_{it} = 1, \\ 0 & \text{if } a_{it} = 0, \end{cases} \quad (4)$$

where $a_{it} = 1$ denotes the decision to work. The participation reward $m(A_t)$ depends on the measure of agents A_t who decide to work, building strategic complementarity. For simplicity $m(A_t)$ is assumed linear in A_t , which may be interpreted as the first order approximation of more involved mechanisms,¹²

$$m(A_t) = \varepsilon + (1 - \varsigma - \varepsilon) A_t,$$

where $\varepsilon, \varsigma > 0$ are arbitrary small numbers ensuring that $m(A_t) \in (0, 1)$.

The cost of participation c_{it} may be interpreted as disutility of working, including the opportunity cost from an alternative activity and a "menu cost." In particular, we interpret this cost as the result of an heterogenous exogenous technology that is additively separable to the complementarity.

Heterogeneity. Agents living at t have heterogeneous costs $c_{it} \in [0, 1]$, represented by

$$f(c_{it} | \theta_t) = \begin{cases} \alpha + \beta & \text{for } c_{it} \in [\theta_t, \theta_t + \sigma], \\ \beta & \text{otherwise,} \end{cases} \quad (5)$$

which is composed of two uniform distributions. There is a density β of agents with cost $c_{it} \in [0, 1]$, but there is a relevant measure of agents with cost inside a smaller range, $c_{it} \in [\theta_t, \theta_t + \sigma]$, with an additional density α . This concentration of agents is called "the cluster," whose relative position in the population depends on θ_t . This parameter is interpreted as the key "fundamental" in this economy since it summarizes the form of the heterogeneity distribution. Thus, higher θ_t implies higher average costs, which we interpret as a lower aggregate productivity. We assume that $\theta_t \in [0, 1 - \sigma]$ to ensure that the cluster is fully contained in the support $[0, 1]$, and restrict parameters α and β to satisfy $F(1 | \theta_t) = 1 \forall \theta_t$, i.e., $\beta = 1 - \alpha\sigma$.

Shocks to the fundamental To make the model dynamic, the next ingredient is a persistent stochastic process for θ_t . To specify this process, we assume that $[0, 1 - \sigma]$ is split into an arbitrarily fine, discrete, and evenly distributed grid $\{w_k\}_{k=1}^K$, such that $w_k = (k - 1)v$, with $k = 1, \dots, K$, $w_1 = 0$, $w_K = 1 - \sigma$, and v denoting the length of each "step" $v = \frac{1 - \sigma}{K - 1}$.

Thus, θ_t follows a similar process to a random walk, although with a couple of twists: It only takes values on the grid $\{w_k\}_{k=1}^K$ and it can move only one step a time with probability $p < \frac{1}{3}$. The

¹²Endogenous motivations for this complementarity are based on, for example, trading frictions (Diamond [13]), demand spillovers (Blanchard and Kiyotaki [6]) or financial constraints (Kiyotaki and Moore [18]).

first assumption ensures that θ_t is bounded in its support. The second assumption imposes that

$$\begin{aligned}\Pr[\theta_{t+1} = w_{k-1} \mid \theta_t = w_k] &= p \text{ if } k > 1; \\ \Pr[\theta_{t+1} = w_{k+1} \mid \theta_t = w_k] &= p \text{ if } k < K.\end{aligned}$$

It says that when θ_t is not located within its boundaries, there is a probability p that θ_{t+1} is located either in the immediate upper or in the immediate lower position of θ_t . Thus, the probability that $\theta_{t+1} = \theta_t$ is $1 - 2p$. However, when θ_t reaches the upper (lower) boundary, there is only probability p that θ_{t+1} is located in the immediate lower (upper) position, and probability $1 - p$ that $\theta_{t+1} = \theta_t$.

This process can alternatively be represented as the vector $\varrho_t = [\varrho_{1t}, \dots, \varrho_{Kt}]'$, such that each of its elements is defined as $\varrho_{kt} = \Pr(\theta_t = w_k)$ for $k = 1, \dots, K$, following a motion described by

$$\varrho_{t+1} = Q\varrho_t \tag{6}$$

where

$$Q = \begin{pmatrix} 1-p & p & & & & \\ p & 1-2p & p & & & \\ & & \ddots & \ddots & \ddots & \\ & & & p & 1-2p & p \\ & & & & p & 1-p \end{pmatrix}$$

This process takes this form to reflect the main message behind Chamley's model: Small shocks may generate occasional large fluctuations. It is straightforward to introduce large shocks, but they make no substantial change in the properties of the equilibrium or in our policy analysis.

Information structure. Unless otherwise specified, agents observe their own participation cost and public information $\mathfrak{S}_t = \{A_{t-1}^*, \mathfrak{S}_{t-1}\}$ containing the history of labor participations. Then, the information set of an agent i living at t is $\{c_{it}, \mathfrak{S}_t\}$.

3.2 Equilibrium multiplicity when the fundamental is observable

Since agents have heterogenous participation costs, the equilibrium is defined as a cutoff strategy in this cost. This section shows that there are multiple equilibrium cutoffs for a range of parameters when θ_t is observable. This is because heterogeneity in costs does not imply heterogeneity in information, so this game is still a coordination game like our starting point in Section 2. The history of past labor participations in this case is irrelevant; thus, agents' decision of working is reduced to a sequence of static decisions. Therefore, we abstract from the use of a time subindex in this analysis. We start by defining the equilibrium under complete information.

Definition 1 *The equilibrium set $\zeta^*(\theta)$ when θ is observable is composed of equilibrium cutoff*

strategies $c^*(\theta)$ such that $a_i = 1$ if and only if $c_i \leq c^*(\theta)$. The equilibrium cutoff $c^*(\theta)$ satisfies the indifference condition of the marginal participant who has $c_i = c^*$:

$$U(1, A(c^*, \theta), c^*) = U(0, A(c^*, \theta), c^*),$$

for a given labor participation $A(c^*, \theta)$ defined by:

$$A(c^*, \theta) = \int_0^{c^*} dF(c_i | \theta).$$

The indifference condition for the marginal participant defining $c^*(\theta)$ is similar to the condition (1) used in global game. However, the main difference is that, due to complete information, agents do not need to form expectations about θ . Strategic complementarity implies that this condition involves a fixed point problem between $c^*(\theta)$ and $A(c^*, \theta)$. To solve this problem, we first compute the indifference condition in this economy for a given equilibrium labor participation $A(c^*, \theta)$:

$$m(A(c^*, \theta)) = \varepsilon + (1 - \varsigma - \varepsilon) A(c^*, \theta) = c^*. \quad (7)$$

And we use the heterogeneity distribution in (5) to compute labor participation $A(c, \theta)$ for an arbitrary cutoff $c \in [0, 1]$ and fundamental θ ,

$$A(c, \theta) = \int_0^c dF(c_i | \theta) = \begin{cases} \beta c & \text{for } c \in [0, \theta]; \\ \beta c + \alpha(c - \theta) & \text{for } c \in (\theta, \theta + \sigma); \\ \alpha\sigma + \beta c & \text{for } c \in [\theta + \sigma, 1]. \end{cases} \quad (8)$$

Hence, we find $c^*(\theta)$ that solves $c^*(\theta) = \varepsilon + (1 - \varsigma - \varepsilon) A(c^*(\theta), \theta)$. Parameters ε and ς in $m(A_t)$ rule out sunspot equilibria with full or zero participation. If the concentration of agents in the cluster is such that $\alpha + \beta > 1$, the equilibrium set $\zeta^*(\theta)$ becomes

$$\zeta^*(\theta) = \begin{cases} \{c_L^*\} & \text{if } \theta > c_H^* - \sigma; \\ \{c_L^*, c_H^*\} & \text{if } \theta \in [c_L^*, c_H^* - \sigma]; \\ \{c_H^*\} & \text{if } \theta < c_L^*; \end{cases} \quad (9)$$

$$\text{with } c_L^* = \frac{\varepsilon}{1 - (1 - \varsigma - \varepsilon)\beta} \text{ and } c_H^* = \frac{\varepsilon + \alpha\sigma(1 - \varsigma - \varepsilon)}{1 - (1 - \varsigma - \varepsilon)\beta}.$$

Figure 1 shows a graphical representation of $m(A(c, \theta))$ of the left-hand side of (7) and a 45° line for its right-hand side for an arbitrary cutoff strategy $c \in [0, 1]$. An equilibrium cutoff c^* is represented by the intersection of these two functions. Given the complementarity ($m' > 0$) and the heterogeneity in costs (5), the function $m(A(c, \theta))$ reminds us of the S-shaped function required for equilibrium multiplicity (Cooper and John [11]).

The equilibrium set (9) and Figure 1A show that two self-fulfilling equilibria exist, c_L^* and c_H^* , when $\theta \in [c_L^*, c_H^* - \sigma]$, i.e., when the cluster is not in the extremes of the distribution. The far left crossing in Figure 1A represents the equilibrium c_L^* , in which agents expect a low labor participation, and because of the externality, they expect a low reward of working that confirms the low participation as an equilibrium. Thus, agents with $c_i \leq c_L^*$ work. Similarly, with the same θ , the far right crossing in Figure 1A represents the equilibrium c_H^* , which is also self-confirmed. Thus, agents with $c_i \leq c_H^*$ work. The fixed point (7) admits another solution, the middle crossing in the Figure 1A, but that solution is ignored in the analysis because it is unstable.¹³

In contrast, there is only one equilibrium when the cluster is in the extremes of the distribution. If $\theta > c_H^* - \sigma$, as in Figure 1B, the cutoff c_L^* is dominant, and if $\theta < c_L^*$, as in Figure 1C, the cutoff c_H^* is dominant. These areas of strict dominance play a key role in producing equilibrium uniqueness when we introduce fundamental uncertainty; the topic to which we now turn.

3.3 Equilibrium uniqueness with fundamental uncertainty

We show now that this economy always has a unique equilibrium when θ_t is unobservable. This is because heterogeneity in costs implies heterogeneity in information, so the Global Games argument in Section 2 also applies here. However, this economy is dynamic and the fundamental is persistent, thus past equilibrium realizations enter into the set of public information. This feature makes the equilibrium cutoff strategy history-dependent, in particular, $c_t^* = c_L^*$ if $c_{t-1}^* = c_L^*$ unless $\theta_t < c_L^*$; and $c_t^* = c_H^*$ if $c_{t-1}^* = c_H^*$ unless $\theta_t > c_H^* - \sigma$. Hence, the equilibrium dynamics of labor participation are characterized by switches between "pessimistic" regimes (when the equilibrium is c_L^* , also called "regime L") and "optimistic" regimes (when the equilibrium is c_H^* , also called "regime H"). We restore the use of a time subindex for this analysis, which we start by defining the equilibrium set under incomplete information.

Definition 2 *The equilibrium set $\zeta^*(\mathfrak{S}_t)$ when θ_t is not observable is composed of equilibrium cutoff strategies $c^*(\mathfrak{S}_t)$ such that $a_{it} = 1$ if and only if $c_{it} \leq c^*(\mathfrak{S}_t)$. The equilibrium cutoff $c^*(\mathfrak{S}_t)$ depends on the public information set $\mathfrak{S}_t = \{A_{t-1}^*, \mathfrak{S}_{t-1}\}$, and must satisfy the indifference condition of the marginal participant who has $c_i = c^*$:*

$$E_{\theta_t} \{U(1, A(c^*, \theta_t), c^*) \mid c^*, \mathfrak{S}_t\} = E_{\theta_t} \{U(0, A(c^*, \theta_t), c^*) \mid c^*, \mathfrak{S}_t\}$$

for a given labor participation $A(c^*, \theta_t)$ defined by:

$$A(c^*, \theta_t) = \int_0^{c^*} dF(c_i \mid \theta_t).$$

¹³ Any disturbance around this solution implies that the equilibrium converges either to c_L^* or c_H^* .

Comparing with Definition 1 in Section 3.2, when θ_t is observable, the indifference condition now holds in expected value. And comparing with the Global Games equilibrium condition (1) in Section 2, exogenous public information z is replaced with the history of labor participations \mathfrak{S}_t .

To solve this problem, we proceed as suggested in Section 2 when payoffs are continuous on labor participation A . This procedure is similar to the one applied when θ_t is observable. The indifference condition for the marginal participant in this economy is

$$E_{\theta_t} [m(A(c^*, \theta_t)) | c^*, \mathfrak{S}_t] = c^* \quad (10)$$

Also, we use $A(c, \theta_t)$ defined in (8) to compute the expected labor participation by the marginal participant. We focus on the case $\alpha + \beta > 1$, which has equilibrium multiplicity when θ_t is observable. The rest of the section is devoted to show that (10) has a unique solution $c^*(\mathfrak{S}_t)$. For illustration, we sketch the argument for a regime L .¹⁴

Regimes. Let us define the "switching thresholds" w_L and w_H as

$$w_L \equiv \sup \{w_k : w_k < c_L^*\}, \quad w_H \equiv \inf \{w_k : w_k > c_H^* - \sigma\}. \quad (11)$$

with c_L^* and c_H^* defined in (9). The term w_L is the highest position of θ_t in its grid $\{w_k\}_{k=1}^K$ such that the cutoff c_H^* is the dominant equilibrium, and w_H is the lowest position of θ_t such that the dominant equilibrium cutoff is c_L^* .

Suppose that agents can observe the fundamental at $t = 0$ in a position $\theta_0 = w_H$ and cannot observe θ_t at any $t > 0$. At $t = 0$ there is almost equilibrium multiplicity, but the unique equilibrium is still c_L^* . At $t = 1$, public information is $\mathfrak{S}_1 = \{\theta_0 = w_H\}$. Note here A_0 is not in the information set because it is redundant information about θ_0 .

Abstracting from private information, the diffusion process of θ_t in (6) implies that the expected participation reward for an arbitrary cutoff strategy c is

$$\begin{aligned} E_{\theta_1} [m(A(c, \theta_1)) | \mathfrak{S}_1] &= pm(A(c, w_{H-1})) + (1 - 2p)m(A(c, w_H)) + pm(A(c, w_{H+1})) \\ &= pm_{-v}(c) + (1 - 2p)m_0(c) + pm_v(c) \equiv m^e(c) \end{aligned}$$

This expression has an explicit solution using the functional form of the reward $m(A)$ and participation $A(c, \theta)$ in (8). Figure 2 shows $m_{-v}(c)$, $m_0(c)$, $m_v(c)$ – the function $m(A(c, \theta_1))$ for θ_1 equal to w_{H-1} , w_H , and w_{H+1} – and a 45° line representing the right-hand side of (10). Thus, an equilibrium is represented by the cross of these curves with the 45° line. If $\theta_1 = w_{H-1}$ is observable, both cutoffs c_L^* and c_H^* are equilibria. However, if either $\theta_1 = w_H$ or $\theta_1 = w_{H+1}$ is observable, c_L^* is the unique equilibrium. Since $m[A(c, \theta_1)]$ is concave around $\theta_1 + \sigma$, the expected payoff $m^e(c)$ is below $m_0(c)$ in the range $[w_{H-1} + \sigma, w_{H+1} + \sigma]$ by Jensen's inequality. Thus, because $m_0(c)$

¹⁴For a formal proof, see Chamley [10].

crosses the 45° line only once, $m^e(c)$ also does the same.

This result implies that when $\theta_0 = w_H$ is observable, c_L^* is the unique equilibrium at $t = 1$ when agents only use public information \mathfrak{S}_1 . This result is reinforced with private information taken into account. This is because agents rationally assign a probability $\frac{\alpha}{\alpha+\beta}$ that they are part of the cluster. Thus, agents with costs in the relevant range $[w_{H-1} + \sigma, w_{H+1} + \sigma]$ assign more probability than p that θ_1 is w_H or w_{H+1} . This flattens the expected participation reward even further, with respect to $m_0(c)$. Hence, the only equilibrium at $t = 1$ is $c_1^* = c_L^*$.

This result can be extended for $t > 1$. That is: c_L^* is persistently selected unless

$$\theta_t < c_L^*, \text{ or equivalently, } \theta_t \leq w_L.$$

Define $\mu_t = [\mu_{1,t}, \dots, \mu_{K,t}]$ as beliefs about the position of θ_t using only public information, with $\mu_{k,t} = \Pr[\theta_t = w_k \mid \mathfrak{S}_t]$ for $k = 1, \dots, K$. Notice that, using (8), $A(c_L^*, \theta_t) = \beta c_L^*$ is constant for any value of θ_t as long as $\theta_t > w_L$. Hence, the ex-post observation of equilibrium c_L^* at t does not reveal the exact position of θ_t . Thus, taking into account the diffusion process of θ_t in (6), beliefs μ_t are updated according to Bayes' rule:

$$\begin{aligned} \tilde{\mu}_{t+1} &= \begin{cases} \frac{\mu_{k,t}}{1 - \sum_{i=1}^L \mu_{i,t}} & \text{for } w_k > w_L, \\ 0 & \text{for } w_k \leq w_L; \end{cases} \\ \mu_{t+1} &= Q\tilde{\mu}_{t+1}, \end{aligned} \tag{12}$$

with μ_0 assigning $\mu_{H,0} = 1$ to the position w_H of θ_t since $\theta_0 = w_H$ is assumed observable. Figure 3A shows a parametrized example of the evolution of the likelihood ϱ_t in (6) and Figure 3B shows beliefs μ_t in (12). The likelihood ϱ_t converges to a uniform $[0, 1]$, while beliefs μ_t skew upwards, converging to

$$\lim_{t \rightarrow \infty} \mu_{k,t} = \begin{cases} 2 \tan\left(\frac{1}{2}r\right) \sin(r(k-L)) & \text{for } w_k > w_L, \\ 0 & \text{for } w_k \leq w_L, \end{cases}$$

where $r = \frac{\pi}{2(K-L)+1}$. This limit function is derived in Appendix A. This skewness ensures that, if c_L^* is the equilibrium at t and if $\theta_t \geq w_L$, c_L^* is the only equilibrium at $t + 1$ when agents only use public information \mathfrak{S}_{t+1} . This result is reflected in Figure 4, where the function m^e crosses the 45° line only in c_L^* . This result is valid when private information is also taken into account, as long as $\alpha < 2(1 - \beta)$ (Chamley [10], proposition 4). The intuition is similar to the case for $t = 1$: Agents assign a high probability to be part of the cluster, placing even greater weight on higher positions of θ_t than in (12) when they update their beliefs μ_t . This effect flattens the expected reward of working even further, reinforcing equilibrium uniqueness.

Expectation switches. Suppose the economy is in a regime L , with equilibrium cutoff c_L^* . Using (8), labor participation is βc_L^* if $\theta_t > w_L$ and $\beta c_L^* + \alpha(c_L^* - \theta_t)$ if $\theta_t \leq w_L$. Hence, if shocks to

the fundamental have accumulated such that $\theta_t \leq w_L$, the ex-post observation of labor participation fully reveals θ_t . In particular, given its slow diffusion process in (6), the information revealed is $\theta_t = w_L$. Therefore, a new regime H with equilibrium cutoff c_H^* arises after applying the same argument as above, only replacing the initial set of public information for $\mathfrak{S}_1 = \{\theta_0 = w_L\}$. This large and sudden change in labor participation after a synchronized revision of expectations is what we call an "expectation switch."

Using the closed form for labor participation $A(c, \theta)$ in (8), participation in the new regime H is $\beta c_H^* + \alpha\sigma$ if $\theta_t < w_H$ and $\beta c_H^* + \alpha(c_H^* - \theta_t)$ if $\theta_t \geq w_H$. Hence, agents cannot infer the exact position of the fundamental if $\theta_t < w_H$, but they can if shocks accumulate such that $\theta_t \geq w_H$. This information put us again in the situation in which a regime L arises, and it closes the cycle.

Therefore, the equilibrium set of cutoff strategies is

$$\zeta^*(\mathfrak{S}_t) = \begin{cases} \{c_L^*\} & \text{if } \mathfrak{S}_t = \{A_{t-1}^*(c_L^*) = \beta c_L^*, \mathfrak{S}_{t-1}\} \\ \{c_H^*\} & \text{if } \mathfrak{S}_t = \{A_{t-1}^*(c_L^*) > \beta c_L^*, \mathfrak{S}_{t-1}\} \\ \{c_H^*\} & \text{if } \mathfrak{S}_t = \{A_{t-1}^*(c_H^*) = \alpha\sigma + \beta c_H^*, \mathfrak{S}_{t-1}\} \\ \{c_L^*\} & \text{if } \mathfrak{S}_t = \{A_{t-1}^*(c_H^*) < \alpha\sigma + \beta c_H^*, \mathfrak{S}_{t-1}\} \end{cases} \quad (13)$$

In contrast to the equilibrium set $\zeta^*(\theta)$ when θ_t is observable in (9), agents' heterogeneity and uncertainty about θ_t provide an equilibrium selection device for $\theta_t \in [c_L^*, c_H^* - \sigma]$. The key distinction of this result, with respect to a static global game, is that the equilibrium is history-dependent, in particular, on last period labor participation A_{t-1}^* . In this sense the equilibrium is Markov, because A_{t-1}^* summarizes all the relevant information contained in the whole history of participations. Hence, aggregate dynamics are unique and follow a regime-switching process. A "pessimistic" regime L delivers a low labor participation $A^*(c_L^*) = \beta c_L^*$ as long as $\theta_t \geq c_L^*$ and a "optimistic" regime H delivers a high labor participation $A^*(c_H^*) = \beta c_H^* + \alpha\sigma$ as long as $\theta_t \leq c_L^* - \sigma$. When shocks to the fundamental accumulate, such that these conditions are not satisfied, precise information about the fundamental is revealed, which triggers an expectation switch that provokes a transition to the alternative regime.

These dynamics are still a coordination failure, since a regime L could take place for $\theta_t \in [c_L^*, c_H^* - \sigma]$ while a Pareto superior regime H can be supported with the same fundamental.¹⁵ This is the main motivation for our policy analysis, which we start in the next section.

¹⁵We abstract from environments in which optimism is inefficient. We argue that this possibility changes the form of optimal policy, but it does not affect our main focus: The study of the channels in which the informational role of policy takes place in economies with self-fulfilling dynamics and expectation switches.

4 Introduction of a government

Our novel analysis starts by introducing a government to our laboratory economy and by studying the effects of policy when θ_t is observable. The key finding is that, as in the example of Section 2, distortionary policy affects the ranges in θ_t with equilibrium multiplicity. In particular, taxes on labor income has control on the cutoffs c_L^* and c_H^* . To pave the way to Section 5, when we study the informational role of policy when θ_t is unobservable, we also produce a reduced form of the government's objective.

4.1 Modifications to the set-up

Players. Private agents remain identical to those in Section 3. The government has an infinite horizon and each period specifies policy $\{\tau_t, g_t\}$, where τ_t is a proportional tax on labor income and g_t is a public good financed by the collected taxes. For the sake of simplicity, public debt is not allowed.¹⁶

Timing. Each time period is broken into three stages:

1. The government specifies taxes to be collected at the end of the current period;
2. Agents in a generation simultaneously decide to work or not to work;
3. Labor participation is observed; the specified tax rate is levied; the public good is provided; and information sets of all players are updated.

Payoffs. Agents' utility in (4) is modified to take into account policy $\{\tau_t, g_t\}$,

$$U(a_{it}, c_{it}, A_t, \tau_t, g_t) = \begin{cases} (1 - \tau_t) m(A_t) - c_{it} + \phi(g_t) & \text{if } a_{it} = 1; \\ \phi(g_t) & \text{if } a_{it} = 0, \end{cases} \quad (14)$$

where $\phi(\cdot)$ is the utility of the public good, with $\phi' > 0$, $\phi'' < 0$, $\phi'(0) = \infty$, and $\phi'(\infty) = 0$.

The government maximizes the present discounted value of welfare across agents and generations,

$$\sum_{t=0}^{\infty} \gamma^t \left[\int_0^1 U(a_{it}, c_{it}, A_t, \tau_t, g_t) dF(c_{it} | \theta_t) \right], \quad (15)$$

where γ is the time discount factor. Because agents live one period, the government's objective is represented as the discounted sum of a sequence of "one-period welfare" (inside the brackets), which is the cross-sectional aggregation of agents' utility in the same generation.

¹⁶This assumption eliminates public debts as a state variable in the policy problem. Thus, abstracting from informational concerns, which is the focus of this paper, the policy problem is static.

Heterogeneity and shocks. Both the heterogeneity in cost $F(c_{it} | \theta_t)$ and the stochastic diffusion process of θ_t remain identical to those in Section 3 as defined by (5) and (6).

Information structure. Agents' information set also remains identical to the one in Section 3, which is composed of their participation cost c_{it} (private information) and public information $\mathfrak{S}_t = \{A_{t-1}^*, \tau_t, \mathfrak{S}_{t-1}\}$. Notice that, given the timing of actions, taxes in the current period are also public information. The government has no control of the fundamental θ_t and has access only to public information \mathfrak{S}_t .

4.2 Equilibrium with taxes when the fundamental is observable

We now study policy when the observation of θ_t may generate two equilibrium cutoffs. As in Section 3.2, we also abstract from a time subindex.

Following Definition 1, an equilibrium cutoff strategy c^* satisfies the indifference condition for the marginal worker. Parallel to Section 2, public goods g represent lump-sum policy, with no effect on the equilibrium because it does not introduce distortions in this indifference condition. Instead, income taxes τ have an effect, such that

$$(1 - \tau) m(A(c^*, \theta)) = c^*, \quad (16)$$

which is the counterpart of (7) in Section 3.2. This equation has explicit solution, such that, when $\alpha + \beta > 1$, the equilibrium set with taxes $\zeta^*(\theta, \tau)$ is

$$\zeta^*(\theta, \tau) = \begin{cases} \{c_L^*(\tau)\} & \text{if } \theta > c_H^*(\tau) - \sigma \\ \{c_L^*(\tau), c_H^*(\tau)\} & \text{if } \theta \in [c_L^*(\tau), c_H^*(\tau) - \sigma] ; \\ \{c_H^*(\tau)\} & \text{if } \theta < c_L^*(\tau) \end{cases} \quad (17)$$

$$\text{with } c_L^*(\tau) = \frac{\varepsilon}{\frac{1}{1-\tau} - (1-\zeta-\varepsilon)\beta} \text{ and } c_H^*(\tau) = \frac{\varepsilon + \alpha\sigma(1-\zeta-\varepsilon)}{\frac{1}{1-\tau} - (1-\zeta-\varepsilon)\beta}.$$

Notice that taxes affect the cutoffs that govern the labor participation in each equilibrium, as well as the range of the fundamental θ with equilibrium multiplicity. The following proposition states some properties of this control for later reference.

Proposition 1 *For $s = L, H$, the equilibrium cutoff strategies c_s^* depend on taxes τ , such that:*

- (i) $\frac{\partial c_s^*}{\partial \tau} < 0$, where the complementarity amplifies the negative effect of taxes on c_s^* ;
- (ii) $\frac{\partial c_H^*}{\partial \tau} < \frac{\partial c_L^*}{\partial \tau}$, taxes have a stronger effect on c_H^* than on c_L^* ;
- (iii) $\frac{\partial^2 c_s^*}{\partial \tau^2} > 0$, taxes have an increasing effect on cutoffs;
- (iv) when the equilibrium is unique on c_H^* , a positive shift of taxes may generate multiplicity

$\{c_L^*, c_H^*\}$ or uniqueness on c_L^* ;

(v) when the equilibrium is unique on c_L^* , a negative shift of taxes may generate multiplicity $\{c_L^*, c_H^*\}$ or uniqueness on c_H^* .

The proof of this proposition is direct from $c_L^*(\tau)$ and $c_H^*(\tau)$ in (17). Property (i) states the negative effect of taxes on the equilibrium cutoffs, and emphasizes the amplification effect generated by the complementarity, just as in the static example of Section 2. The intuition behind this amplification effect becomes clear after differentiating the equilibrium condition (16),

$$\frac{dc_s^*(\tau)}{d\tau} = -m[A(c^*, \theta)] + (1 - \tau) \frac{\partial m}{\partial A} \frac{\partial A}{\partial c_s^*} \frac{dc_s^*}{d\tau} < 0 \quad (18)$$

A higher tax decreases incentives to work, which in turn decreases the cutoff in either equilibrium. This effect is captured by the first term on the right-hand side of (18). In addition, because the complementarity in $m[\cdot]$, lower participation feeds back on even lower incentives to work, amplifying the negative effect of tax increases on cutoffs. This feedback is captured by the second term on the right-hand side of (18). If there is no complementarity, $\frac{\partial m}{\partial A} = 0$, there is no such amplification effect.

The rest of this proposition presents other properties of this control. Property (ii) states that there is a bolder effect of taxes on c_H^* than on c_L^* because taxes are proportional and the reward of working is higher for c_H^* than for c_L^* . Property (iii) remarks that, with a similar mechanism as in property (i), taxes have an increasingly negative effect on equilibrium cutoffs. Property (iv) stresses that, when a higher tax decreases $c_L^*(\tau)$, it expands the range of θ where either c_L^* is strictly dominant or there are multiple equilibria. Thus it is more likely for the fundamental θ to fall into that range. Property (v) states the mirror image of property (iv). A lower tax increases $c_H^*(\tau)$ and expands the range of θ where either c_H^* is strictly dominant or there are multiple equilibria.

4.3 Reduced form of the policy objective

We do not study optimal policy when θ_t is observable. However, we use the equilibrium set $\zeta^*(\theta, \tau)$ to obtain a reduced form of the policy objective (15) that will be used in the next section.

The government's problem may be represented in the familiar form of an optimal control problem that maximizes the government's objective (15) subject to the competitive equilibrium denoted by $\zeta^*(\theta_t, \tau_t)$ in (17),

$$\max_{\{\tau_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \gamma^t \left[\int_0^1 U(a, c_i, A_t, \tau_t, g_t) dF(c_i | \theta_t) \right] \text{ s.t. } \zeta^*(\theta_t, \tau_t).$$

Since there is no state variable in agents' decisions, such as capital and public debts, there is no need for using the Ramsey's approach (as in Chamley [9]) or commitment multipliers (as in

Marcet and Marimon [21]). Thus, we concentrate our attention on finding a reduced form for the one-period welfare

$$\int_0^1 U(a, c_i, A_t, \tau_t, g_t) dF(c_i | \theta_t).$$

To start, since only agents with $c_i \leq c_s^*$ work in equilibrium $s = L, H$, one-period welfare may be represented as

$$\int_0^{c_s^*} \{(1 - \tau_t) m(A_t) - c_i\} dF(c_i | \theta_t) + \phi(g_t).$$

This expression reflects that, while the reward of not working is zero, all agents enjoy the public good. Furthermore, all workers receive the same reward but have heterogeneous costs of participation. Thus, one-period welfare may be further simplified,

$$(1 - \tau_t) m(A(c_s^*, \theta)) A(c_s^*, \theta) - \bar{E}(c_i \leq c_s^* | \theta_t) + \phi(g(c_s^*, \tau_t)),$$

for equilibrium $s = L, H$. The first term is the aggregate consumption, which is the after-tax reward of working multiplied by the number of participants. The second one is the cross-sectional sum of costs for those who work. And the third term is the utility generated by the public good.

The equilibrium condition (16) is useful for further simplification of the one-period welfare,

$$c_s^* A(c_s^*, \theta) - \bar{E}(c \leq c_s^* | \theta_t) + \phi(g(c_s^*, \tau_t)) \equiv Z(c_s^*, \theta_t, \tau_t). \quad (19)$$

There is a closed form solution for each of these terms. The equilibrium cutoffs $c_s^*(\tau_t)$ are obtained in (17). Labor participation $A(c_s^*)$ is obtained by using (8) and the fact that, if c_L^* takes place, the entire cluster does not work, and if c_H^* takes place, the entire cluster works:

$$A(c_s^*) = \beta c_s^*(\tau_t) + \mathbf{1}_{s=H} \cdot \alpha \sigma, \quad (20)$$

where $\mathbf{1}_{s=H} = 1$ if the equilibrium cutoff is $c_H^*(\tau_t)$. The average participation cost is obtained using the cost heterogeneity function $F(c_i | \theta_t)$:

$$\bar{E}(c_i \leq c_s^* | \theta_t) = \frac{\beta}{2} (c_s^*)^2 + \mathbf{1}_{s=H} \left[\alpha \sigma \left(\theta_t + \frac{1}{2} \sigma \right) \right], \quad (21)$$

and the assumption of no public debt implies that

$$g_s(c_s^*, \tau_t) = \frac{\tau_t}{1 - \tau_t} c_s^* A(c_s^*). \quad (22)$$

Therefore, the government's objective when θ_t is observable may be represented as

$$\max_{\{\tau_t\}_{t=0}^{\infty}} \left\{ \sum_{t=0}^{\infty} \gamma^t Z(c_t^*, \theta_t, \tau_t) \right\} \text{ s.t. } c_t \in \zeta^*(\theta_t, \tau_t). \quad (23)$$

One-period welfare $Z(\cdot)$ critically depends on which equilibrium cutoff takes place, c_L^* or c_H^* . The cutoff determines labor participation, reward of working, average costs, and the provision of the public good. The fundamental θ_t and taxes τ_t also affect one-period welfare, but only with a minor direct effect: θ_t enters in the average cost of participation, and τ_t enters in the provision of the public good. However, both the fundamental θ_t and taxes τ_t have an important indirect effect: the ranges with equilibrium multiplicity are defined on θ_t ; and τ_t controls the level of the cutoffs, $c_L^*(\tau_t)$ or $c_H^*(\tau_t)$. This control of taxes will be the driving force of our results in the next section.

5 Optimal policy under fundamental uncertainty

We now focus on the central contribution of the paper: The ways in which the control of taxes on $c_L^*(\tau_t)$ and $c_H^*(\tau_t)$ translates into an *informational role* of policy when the fundamental θ_t is unobservable, i.e., when there is a unique equilibrium with regime-switching dynamics in our economy, as in Section 3.3. As in the global game of Section 2, this control of taxes affects welfare through four channels: (a) on one-period welfare, i.e., aggregate consumption, average participation costs and public goods provision in either regime; (b) on equilibrium selection, i.e., the probability of expectation switches; (c) on learning about θ_t , i.e., future expectation switches and on the expected duration of regimes; and (d) on recovering equilibrium multiplicity, i.e., provoking non-fundamental volatility.

We first abstract from the channel (d) by imposing the assumption that equilibrium uniqueness is preserved for any sequence of taxes. By doing so, we benefit from recursive representation of the policy problem that exploits the regime-switching dynamics, obtaining a number of results. Policy has little leverage on expectations at the onset of regimes because agents have precise information that a new switch cannot occur in the near term. Hence, policy should concentrate on its other objectives, which in our case is to attain the optimal provision of public goods for each generation.

However, as time passes in a regime and information becomes less precise, there are higher incentives for deviating from this objective to implement a stabilization policy. When taxes are cut, switching probability increases in a pessimistic regime L and decreases in an optimistic regime H —the channel (b). Because a regime H is socially more desirable than a regime L , tax cuts are implemented in spite of their costs in welfare of the current generation—the channel (a)—and its effects on agents’ learning about fundamentals – the channel (c). This channel (c) operates differently in a regime L than in a regime H , introducing asymmetry in the design of taxation schemes. Severe and transitory tax cuts ("big pushes") are optimal attempts to break pessimism, while small and permanent tax cuts are optimal attempts to extend optimism.

We study the channel (d) after relaxing the assumption that policy has no effect on equilibrium uniqueness. We find that a successful big push breaking pessimism may generate equilibrium multiplicity. However, we propose a fix. A tax scheme contingent on which equilibrium takes place may rule out this source of multiplicity. This policy design theoretically may eliminate the low

participation equilibrium, but it is only feasible at the onset of the optimistic regime.

5.1 Optimal policy problem assuming equilibrium uniqueness

This section displays our analysis under the assumption of unique equilibrium dynamics.

5.1.1 Recursive representation

The reduced form of the government's objective (23) has the following recursive representation:

$$W(\mu_t) = \max_{\tau_t \in [0,1]} \{E_t [Z(c_L^*, \theta_t, \tau_t)] + \gamma(1 - \pi_t(c_L^*)) W(\mu_{t+1}) + \gamma\pi_t(c_L^*) V(\mu_1)\}, \quad (24)$$

$$V(\mu_t) = \max_{\tau_t \in [0,1]} \{E_t [Z(c_H^*, \theta_t, \tau_t)] + \gamma(1 - \pi_t(c_H^*)) V(\mu_{t+1}) + \gamma\pi_t(c_H^*) W(\mu_1)\}. \quad (25)$$

$W(\mu_t)$ and $V(\mu_t)$ denote the welfare in the regime $s = L, H$, with a state vector μ_t representing the government's beliefs about the position of θ_t in its grid $\{w_k\}_{k=1}^K$. The control variable is the tax rate $\tau_t \in [0, 1]$. The momentary objective $Z(c_s^*, \theta_t, \tau_t)$ represents one-period welfare in (19). Since θ_t is unobservable, the government uses its beliefs μ_t to compute expectations E_t on this term. The subindex t represents time elapsed from the last switch and γ is the time discounting.

Given beliefs μ_t , the government computes the probability that a regime survives, $1 - \pi_t(c_s^*)$. If this is the case, the continuation value is represented by $W(\mu_{t+1})$ and $V(\mu_{t+1})$, with government's updated beliefs μ_{t+1} . Conversely, if there is an expectation switch at the end of t , the continuation value is $W(\mu_1)$ and $V(\mu_1)$, with government's beliefs at the onset of the alternative regime, μ_1 . The subindex 1 denotes one period after the switch.

With this recursive representation at hand, we show the policy trade-offs involved by studying the effect of policy on each component of (24) and (25).

5.1.2 Static effects

We first study the channel (a) of the informational role of policy, i.e., the effect of taxes on expected one-period welfare in (24) and (25). Based on (19), the expected one-period welfare is defined as:¹⁷

$$E_t [Z(c_s^*, \theta_t, \tau_t)] = c_s^*(\tau_t) A(c_s^*(\tau_t)) - \bar{E}_t [c_{i,t} | c_{i,t} \leq c_s^*(\tau_t)] + \phi(g_s(\tau_t)), \quad (R1)$$

where the restriction of no public debt implies that

$$g_s(\tau_t) = \frac{\tau_t}{1 - \tau_t} c_s^*(\tau_t) A(c_s^*(\tau_t)), \quad (R2)$$

¹⁷The only difference caused by taking expectations is the term of average participation cost. Because the expected participation—in turn the terms of aggregate consumption and public good provision—depends only on the equilibrium cutoffs (equation (20)), rather than on θ_t .

These terms, (R1) and (R2), enter as restrictions in the problem (24) and (25). Recall that, according to Proposition 1, taxes have a negative effect on cutoffs $c_s^*(\tau_t)$. This effect is amplified by the strategic complementarity that arises from the positive effect of labor participation A_t on the reward of working $m(A_t)$. Also recall that one-period welfare critically depends on these cutoffs.

Thus, taxes have a negative effect on the first term of (R1), which is aggregate consumption. This is because a lower cutoff $c_s^*(\tau_t)$ decreases the equilibrium reward of working, which coincides with $c_s^*(\tau_t)$, and decrease labor participation, which also depends on $c_s^*(\tau_t)$ according to (20). In addition, taxes have a positive effect on the second term of (R1) because a lower cutoff $c_s^*(\tau_t)$ decreases the average cost of participation, according to (21). Finally, taxes have an ambiguous effect on the third term of (R1) because a lower $c_s^*(\tau_t)$ could imply either higher or lower tax revenues, according to the standard Laffer curve implicit in (R2).

The next definition establishes a benchmark tax rate for later analysis.

Definition 3 *The static optimal tax $\bar{\tau}_s^*$ is defined for a regime $s = L, H$ as*

$$\bar{\tau}_s^* = \arg \max_{\tau_t \in [0,1]} E_t [Z(c_s^*, \theta_t, \tau_t)]$$

This tax maximizes expected one-period welfare, where the concavity of $\phi(\cdot)$ ensures interior solution. This tax rate is time-invariant in a given regime $s = L, H$, and represents the policy of a benevolent government that ignores the intertemporal effects of the informational role of policy and thus focuses only on its short-run objectives.

5.1.3 Effects on switching probability

Section 3.3 shows that an expectation switch occurs when θ_t reaches its switching threshold, either w_L or w_H . This is because, in this situation, agents can infer the exact position of θ_t by the ex-post observation of the labor participation. The switching thresholds depend on equilibrium cutoffs, which now depend on taxes, $c_L(\tau_t)$ and $c_H(\tau_t)$. This effect is equivalent to the channel (b) in the static example of Section 2.

An expectation switch is triggered when $\theta_t < c_L(\tau_t)$ if the economy is in regime L , or when $\theta_t > c_H(\tau_t) - \sigma$ if the economy is in regime H . Thus, the switching thresholds are now defined as $w_{L(\tau_t)}$ and $w_{H(\tau_t)}$:

$$w_{L(\tau_t)} \equiv \sup \{w_k : w_k < c_L^*(\tau_t)\}, \quad w_{H(\tau_t)} \equiv \inf \{w_k : w_k > c_H^*(\tau_t) - \sigma\}. \quad (26)$$

Thus, given beliefs μ_t , the government computes the switching probability:

$$\pi_t(c_s^*(\tau_t)) = \begin{cases} \Pr[\theta_t < c_L(\tau_t) \mid \mu_t] = \sum_{j=1}^{L(\tau_t)} \mu_{j,t} & \text{for } s = L, \\ \Pr[\theta_t > c_H(\tau_t) - \sigma \mid \mu_t] = \sum_{j=H(\tau_t)}^K \mu_{j,t} & \text{for } s = H, \end{cases} \quad (R3)$$

where $L(\tau_t)$ and $H(\tau_t)$ are the respective position of $w_{L(\tau_t)}$ and $w_{H(\tau_t)}$ on the grid $\{w_k\}_{k=1}^K$. This term (R3) also enters as a restriction in the policy problem (24) and (25). Because of the negative effect of taxes on cutoffs, higher taxes decrease both $w_{L(\tau_t)}$ and $w_{H(\tau_t)}$.¹⁸

Decreasing switching thresholds $w_{L(\tau_t)}$ and $w_{H(\tau_t)}$ have opposite effects on switching probabilities among regimes, however, both have a negative effect on welfare. According to (R3), a lower $w_{L(\tau_t)}$ implies a smaller number of positions of θ_t that can trigger a switch from pessimism to optimism. Thus, higher taxes decrease switching probability $\pi_t(c_L^*(\tau_t))$ in a regime L . Recall that the regime H implies, for the same fundamental, higher aggregate consumption and higher public good provision. Therefore, a lower switching probability in a regime L delivers lower expected welfare. In the opposite case, the regime H , lower $w_{H(\tau_t)}$ implies more positions of θ_t that can trigger a switch to pessimism. Thus, higher taxes increase switching probability $\pi_t(c_H^*(\tau_t))$ in a regime H and have a negative effect on welfare, too. In the case where taxes are cut, there is a symmetric contrary effect.

5.1.4 Effects on learning

Regardless of whether or not there is an expectation switch, the updating of agents' (and the government's) beliefs about θ_t depend on the switching thresholds $w_{L(\tau_t)}$ and $w_{H(\tau_t)}$, which in turn are dependent on taxes, according to (26). This effect is equivalent to the channel (c) in the example of Section 2.

When there is no expectation switch at the end of $t - 1$, i.e., the regime s continues, beliefs evolve according to Bayes' rule as in (12):

$$\tilde{\mu}_t = \begin{cases} \frac{\mu_{kt-1}}{1-\pi_t(c_s^*)} & \begin{cases} \text{for } w_k > w_{L(\tau_{t-1})} \text{ (regime } s = L), \\ \text{for } w_k < w_{H(\tau_{t-1})} \text{ (regime } s = H); \end{cases} \\ 0 & \text{otherwise;} \end{cases} \quad (\text{R4})$$

where $\tilde{\mu}_t$ denotes updated beliefs, with $\tilde{\mu}_{kt} = \Pr[\theta_{t-1} = w_k \mid \mathfrak{S}_t]$ for $k = 1, \dots, K$, and \mathfrak{S}_t is the set of public information, $\mathfrak{S}_t = \{A_{t-1}^*, \tau_t, \mathfrak{S}_{t-1}\}$. Hence, the truncation point imposed by Bayes' rule in (R4) depends on taxes.

In the contrary case, when an expectation switch occurs at the end of $t - 1$, θ_{t-1} is ex-post revealed to be at the position $w_{L(\tau_{t-1})}$ and $w_{H(\tau_{t-1})}$, respectively. Thus, beliefs are updates as

$$\tilde{\mu}_{k1} = \begin{cases} 1 & \text{for } w_k = w_{s(\tau_{t-1})}, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{R5})$$

for regime $s = L, H$. Hence, what agents learn during an expectation switch also depends on taxes.

¹⁸The discretization of the support of θ_t implies a discontinuity in the effect of taxes on these switching thresholds. We abstract from this discontinuity in the following sections, such that we refer as a tax change to a shift large enough to change these thresholds.

Finally, the diffusion process (6) of θ_t implies that

$$\mu_t = Q\tilde{\mu}_t \quad \forall t \geq 1. \quad (\text{R6})$$

These terms (R4), (R5) and (R6) characterize the evolution of the state vector μ_t , and also enter as restrictions in the policy problem (24) and (25).

Table 1 summarizes all relevant cases of the effect of taxes on welfare through this channel.

Tax cut:	Switch	No switch
Regime L	expected shorter upcoming optimism, negative effect on welfare	lower future switching prob., negative effect on welfare
Regime H	switching probability is zero, with positive effect on welfare; but higher future switching prob.	
Tax increase:		
Regime L	switching probability is zero, with negative effect on welfare; but higher future switching prob.	
Regime H	expected shorter upcoming pessimism, positive effect on welfare	lower future switching prob., positive effect on welfare

Table 1 – Effect of taxes on information extraction and learning
in regimes with low and high participation

The first row focuses on a regime L when taxes are cut. The relevant switching threshold in a regime L is $w_{L(\tau_t)}$, which is in the lower part of the support of the distribution μ_t denoting beliefs about θ_t . As shown in Section 4.2.3, the switching probability increases in this regime if taxes are cut. If an expectation switch takes place, the information revealed is that $\theta_t = w_{L(\tau_t)} > w_{L(\tau_{t-1})}$. The distance between this starting point $w_{L(\tau_t)}$ and the new switching threshold $w_{H(\tau)}$ for the new regime H is shorter for any τ with respect to the case in which the starting point is $w_{L(\tau_{t-1})}$. Hence, the diffusion of beliefs in (R6) implies that switching probability in the new regime H increases faster. The tax cut thus has a negative marginal effect on welfare through this channel, conditional on the event of an expectation switch.

If no expectation switch occurs in a regime L after taxes are cut, the fundamental is revealed to be worse than previously expected ($\theta_t > w_{L(\tau_t)} > w_{L(\tau_{t-1})}$ since $\tau_t < \tau_{t-1}$). Hence, expectations are revised down, increasing the expected duration of the current pessimism. In this case, the tax cut therefore has a negative marginal effect on welfare through its effects on learning.

We now turn to tax cuts in a regime H . The relevant switching threshold in this regime is $w_{H(\tau_t)}$, which is in the higher part of the support of the distribution μ_t denoting beliefs about θ_t . As shown in Section 4.2.3, the switching probability decreases in this regime if taxes are cut. This effect is enhanced through the learning channel, turning the switching probability at t to zero. This

is because, after no expectation switch at $t - 1$, all agents infer that $\theta_{t-1} < w_{H(\tau_{t-1})} < w_{H(\tau_t)}$ since $\tau_t < \tau_{t-1}$. Hence, the slow diffusion of θ_t implies that $\theta_t \leq w_{H(\tau_t)}$, i.e., the switching probability in the current period π_t is zero. Notice that in this case the truncation implied by Bayes' rule in (R4) has no effect on beliefs μ_{t+1} and only (R6) applies. This implies higher future switching probability, which conveys a negative side effect on welfare.

The lower part of Table 1, when taxes are raised, is the mirror image of the upper part, when taxes are cut. Hence, a tax increase turns switching probability to zero in a regime L . Despite the fact that a tax increase implies higher switching probability in a regime H , it has a positive side effect through the learning channel, regardless of whether there is an expectation switch or not.

5.2 Optimal policy assuming equilibrium uniqueness

This section studies the optimal policy problem (24) and (25) subject to restrictions (R1) to (R6). The efficiency conditions for the government's problem in each regime are

$$\frac{\partial E_t [Z_t(c_L^*, \theta_t, \tau_t)]}{\partial \tau_t} + \gamma \frac{\partial c_L^*}{\partial \tau_t} \left\{ \frac{\partial \pi_t}{\partial c_L^*} [V_1(\cdot) - W_{t+1}(\cdot)] + (1 - \pi_t) \frac{\partial W_{t+1}(\cdot)}{\partial c_L^*} + \pi_t \frac{\partial V_1(\cdot)}{\partial c_L^*} \right\} = 0 \quad (27)$$

$$\frac{\partial E_t [Z(c_H^*, \theta_t, \tau_t)]}{\partial \tau_t} + \gamma \frac{\partial c_H^*}{\partial \tau_t} \left\{ -\frac{\partial \pi_t}{\partial c_H^*} [V_{t+1}(\cdot) - W_1(\cdot)] + (1 - \pi_t) \frac{\partial V_{t+1}(\cdot)}{\partial c_H^*} + \pi_t \frac{\partial W_1(\cdot)}{\partial c_H^*} \right\} = 0 \quad (28)$$

The notation is simplified in these expressions, such that $V_t(\cdot) \equiv V(\mu_t)$ and $W_t(\cdot) \equiv W(\mu_t)$. The first term represents the *static effects* on aggregate consumption, participation costs, and the public good provision in a given regime. The remaining terms represent the dynamic effect of policy, which operates solely through its control on cutoffs $c_s^*(\tau_t)$, with $s = L, H$. The first term in the parenthesis captures the effect of policy on welfare through *switching probability*, which is weighted by the difference in the continuation values of optimistic and pessimistic regimes. The last two terms summarize the effect of taxes on future welfare through *learning* or the evolution of beliefs, which depends on whether the same regime continues with probability $1 - \pi_t$ or there is a switch with probability π_t .

5.2.1 Basic properties

We now turn to illustrate some basic properties of the optimal tax scheme arising from (27) and (28). Although our policy problem has no closed form solution, these properties help to characterize its time path along the business cycle. These properties are also exploited in the appendix B, where this problem is numerically solved for a parametrized example.

We focus on the case where the difference in participation between optimistic and pessimistic regimes is relevant enough to ever motivate a stabilization policy. Specifically, we restrict param-

ters of both the heterogeneity in costs (5) and the diffusion process (6) of θ_t such that

$$w_{H(\bar{\tau}_H^*)} - w_{L(0)} \geq \kappa \quad (29)$$

for some $\kappa > 1$, so that $c_L^*(0) < c_H^*(\bar{\tau}_H^*) - \sigma$, with $\bar{\tau}_H^*$ specified by Definition 3.

Proposition 2 below shows the first of these properties, in which the continuous policy control $\tau_t \in [0, 1]$ can be reduced to a discrete set τ^s ($s = L, H$).

Proposition 2 *The set of policy options τ^s in the regime $s = L, H$ is finite, with elements τ_k^s ,*

$$\tau_k^s = \begin{cases} \sup \{ \tau \in [0, 1] : w_{s(\tau)} = w_k \} & \text{if } w_k < w_{s(\bar{\tau}_s^*)}, \\ \bar{\tau}_s^* & \text{if } w_k = w_{s(\bar{\tau}_s^*)}, \\ \inf \{ \tau \in [0, 1] : w_{s(\tau)} = w_k \} & \text{if } w_k > w_{s(\bar{\tau}_s^*)}, \end{cases}$$

for $k = 1, \dots, K$, and $\text{size}(\tau^L) < \text{size}(\tau^H) \leq K$.

Proof. Define the sets $\Gamma_k^s = \{ \tau \in [0, 1] : w_{s(\tau)} = w_k \}$ for $s = L, H$. Any $\tau \in \Gamma_k^s$ delivers the same switching probability π_t in (R3), and the same evolution of beliefs μ_{t+1} in (R4), (R5), and (R6). Hence, the only difference between tax rates $\tau \in \Gamma_k^s$ is the one-period welfare, which is concave in τ . By definition, the rate $\bar{\tau}_s^*$ maximizes one-period welfare, so the tax rate $\tau \in \Gamma_k^s$ with the highest welfare is the $\sup \Gamma_k^s$ if $w_k < w_{s(\bar{\tau}_s^*)}$; the $\inf \Gamma_k^s$ if $w_k > w_{s(\bar{\tau}_s^*)}$; and $\bar{\tau}_s^*$ if $w_k = w_{s(\bar{\tau}_s^*)}$, for $s = L, H$. To finish, notice that if some w_k is not reachable for $\tau \in [0, 1]$, $\tilde{\tau}_k^s = \{\emptyset\}$. Thus, using Proposition 1, which states that $\frac{\partial c_H^*(\tau_t)}{\partial \tau_t} < \frac{\partial c_L^*(\tau_t)}{\partial \tau_t} < 0$, this implies that $\text{size}(\tau^L) < \text{size}(\tau^H) \leq K$. ■

This proposition simplifies the policy problem from choosing a sequence of tax rates in a regime $s = L, H$ to choosing a sequence of reachable thresholds $w_{s(\tau_t)}$, and then recovering taxes using τ^s .

We now obtain the optimal tax at the onset of a regime, i.e., one period after a switch.

Proposition 3 *At the onset of a regime L , the optimal tax rate is $\tau_{1,L}^* = \bar{\tau}_L^*$ if the tax rate levied in the last period of the previous regime H is $\tau_{t,H} \leq \bar{\tau}_H^*$.*

At the onset of a regime H , the optimal tax rate is $\tau_{1,H}^ = \bar{\tau}_H^*$ for any tax rate levied in the last period of the previous regime L .*

Proof. *Regime L .* Beliefs μ_1 after a switch assign positive probabilities to positions $w_{H(\tau_t)-1}$, $w_{H(\tau_t)}$, and $w_{H(\tau_t)+1}$ according to (R5) and (R6), where τ_t is the tax levied in the last period before the switch. After imposing that $\tau_{t,H} \leq \bar{\tau}_H^*$, $w_{H(\tau_t)-1} > w_{L(0)}$ by condition (29), since $w_{H(\bar{\tau}_H^*)-1}$ is the lowest possible switching threshold in a regime H and $w_{L(0)}$ is the highest possible switching threshold in the new regime L .¹⁹ Hence $\pi_1(\tau_1) = 0 \forall \tau_1 \in \tau^L$ using (R3), with τ^L defined in Proposition 2. Thus, there are no intertemporal effects of policy, so $\tau_1^* = \bar{\tau}_L^*$.

¹⁹Proposition 6 shows that the restriction $\tau_{t,H} \leq \bar{\tau}_H^*$ is satisfied in equilibrium.

Regime H. Similarly, beliefs μ_1 after a switch at t assign positive probabilities to positions $w_{L(\tau_t)-1}$, $w_{L(\tau_t)}$, and $w_{L(\tau_t)+1}$, where $w_{L(\tau_t)+1}$ is lower than $w_{H(\bar{\tau}_H^*)}$, the new relevant switching threshold. By condition (29), $\pi_1(\tau_1) = 0 \forall \tau_1 \leq \bar{\tau}_H^*$, so $\bar{\tau}_H^*$ maximizes welfare in this range. A rate $\tau_1 > \bar{\tau}_H^*$ may generate an undesirable switch, which lowers the welfare relative to the case with tax rate $\bar{\tau}_H^*$. Thus, $\tau_1^* = \bar{\tau}_H^*$. ■

This proposition pines down the starting point of the sequence of optimal taxes in both regimes based on the fact that information about the fundamental is quite precise at a position far from the new switching threshold. This allows us to use $\bar{\tau}_L^*$ and $\bar{\tau}_H^*$ as benchmarks to evaluate the relevance of the dynamic effects of the informational role of our taxation policy.

5.2.2 Optimal policy under pessimism (regime L)

We now turn to characterize the sequence of optimal taxes in the regime L for $t > 1$. Based on Section 5.1, there are a number of trade-offs that the government faces in this regime for its policy options τ^L . Any tax increase, $\tau_{t,L} > \tau_{t-1,L}$ turns $\pi_t = 0$ by (R3) and (R4), with a negative effect on expected welfare. Decreasing taxes, $\tau_{t,L} < \tau_{t-1,L}$, however, increases π_t by (R3), improving welfare by that channel, but with negative effects through one-period welfare (R1) if $\tau_{t,L} < \tau_{t-1,L} \leq \bar{\tau}_L^*$, and through learning. Regarding the learning effect, if an expectation switch does not occur, probabilities of future switches are revised down by (R4) and (R6); if an expectation switch does occur, the expected duration of the upcoming regime H is lower than the case where the switch is triggered with higher taxes, by (R5) and (R6). Another possibility is to keep taxes constant, $\tau_{t,L} = \tau_{t-1,L}$, but this policy could inefficiently delay the expectation switch. The next proposition is the core result of the paper for policy in "pessimistic" times, putting in perspective the order of magnitude of these effects.

Proposition 4 *When (29) is satisfied for some $\kappa \geq \tilde{\kappa}$, there are only two admissible policy options $\forall t$ in a regime L : (i) $\tau_{t,L}^* = \bar{\tau}_L^*$; or (ii) $\tau_{t,L}^* = \inf \{\tau^L\}$, which is called "a big push."*

This proposition is proved in Appendix A. This result simplifies the policy problem during a regime L , reducing it to a sequence of binomial decisions. Proposition 3 shows that the optimal tax rate at the onset of the regime is the static-optimal rate $\bar{\tau}_L^*$. Meanwhile, the switching probability increases with time as μ_t spreads out over the support of θ_t , according to (R4) and (R6). Hence, when the switching probability reaches some particular level, the government significantly decreases taxes: a big push. The condition $\kappa \geq \tilde{\kappa}$ is useful to pin down a specific rule for these big pushes, $\tau_{t,L}^* = \inf \{\tau^L\}$, but in qualitative terms this condition is not necessary. The key property is that, if the government deviates from its short run objectives (i.e. if $\tau_{t,L}^* \neq \bar{\tau}_L^*$), it has strong incentives to implement large tax cuts because its marginal gain in switching probability increases with the magnitude of the cut. As we show below, this feature is not present in "optimistic" times.

The next proposition shows the optimal policy if a big push fails to generate a switch.

Proposition 5 *If a big push implemented at t fails to trigger a switch, incentives to keep the push at $t + 1$ strongly decreases, so in most cases $\tau_{t+1,L}^* = \bar{\tau}_L^*$.*

Proof. After the failure of a big push, the government (and private agents) infers that $\theta_t > w(\tau_{t,L}^*) > w(\bar{\tau}_L^*)$. Beliefs are updated using this information according to the Bayes' rule, as in (R4) and (R6), implying $\pi_{t+1}(c_L^*(\tau)) = 0$ for all $\tau > \tau_{t,L}^*$. Such a truncation at the left tail of the belief function μ_{t+1} substantially reduces the switching probability in period $t + 1$ if the big push is kept. On the other hand, the one-period welfare loss involved in the big push remains as large as in period t . In most cases where the reduction in switching probability is important, the cost-benefit analysis suggests to raise taxes. Because any rate higher than $\tau_{t,L}^*$ generates zero switching probability, i.e., few intertemporal effects, the best rate to return to is then the static optimal one $\tau_{t+1,L}^* = \bar{\tau}_L^*$. ■

In summary, the benchmark tax $\bar{\tau}_L^*$ is optimal at the onset of a pessimistic regime. As time elapses in this regime, a policy of oscillatory taxes is implemented in which a big push is applied once in a while. If the push fails, then $\bar{\tau}_L^*$ is levied for some time, until beliefs about θ_t assign a high enough probability of a successful new push.

5.2.3 Optimal policy under optimism (regime H)

As Table 1 shows, the effect of taxes on learning in a regime H operates as the mirror image of its effects in a regime L . Thus, policy trade-offs are also a mirror image of those trade-offs in a regime L . A tax cut, $\tau_{t,H} < \tau_{t-1,H}$ turns expected switching probability $\pi_t = 0$ by (R3) and (R4), with a positive effect on welfare. But such policy produces negative effects because of lower one-period welfare in (R1) if the previous tax rate was already smaller than $\bar{\tau}_H^*$, and because expectations of a future switch increases faster by (R4) and (R6). The contrary policy, raising taxes $\tau_{t,H} > \tau_{t-1,H}$, exposes the economy to a higher switching probability. However, it has positive side effects because 1) it increases one-period welfare if $\tau_{t-1,H} < \bar{\tau}_H^*$; and 2) it has positive effects through learning, regardless of whether there is a switch (shorter expected duration of the upcoming regime L) or not a switch (longer duration of the current regime H).

The effects of policy on learning in a regime H open the door to policy experimentation. For instance, suppose that the government is in the middle of a policy involving a sequence of gradual tax cuts. In this case, the policy maker may want to conduct an "experiment" by increasing taxes for a short period of time. If there is no switch, it learns that θ_t is low enough, so there is no need to support optimism by tax cuts. Hence, it can safely return to a higher tax rate, closer to $\bar{\tau}_H^*$, and delay future tax cuts. A more unconventional policy could even involve a big increase in taxes by the authority, so as to push the economy to a recession but with a minimized expected duration.

The next proposition states the core result of the paper for policy in the regime H .

Proposition 6 *When (29) is satisfied for some $\kappa \geq \tilde{\kappa}$, there are only two policy options $\forall t$ in a regime H : (i) $\tau_{t,H}^* = \tau_{t-1,H}^*$; or (ii) decreasing taxes by one "step," $\tau_{t,H}^* = \sup \left\{ \tau \in \tau^H : \tau < \tau_{t-1,H}^* \right\}$.*

This proposition is also proved in the Appendix. As in Proposition 5, this proposition also simplifies the taxation problem, this time in a regime H , transforming it into a sequence of binary options. This result is built on two key features. Because of Bayesian learning, any tax cut turns $\pi_t = 0$, so a small tax cut has the same intertemporal effect as any larger tax cut. In addition, any form of experimentation is suboptimal, because the endogeneity of switching probability to policy implies that any tax increase exposes the economy to a higher risk of an inefficient adjustment in expectations. If the difference in participation between the two regimes is large enough (i.e., if pessimistic episodes involve relevant costs in welfare), this force rules out experimentation. Due to its importance, this result is highlighted in a corollary.

Corollary 1 *Since equilibrium selection is endogenous to policy, incentives for policy experimentation are smaller than when economic dynamics follow an exogenous regime switching process. Thus, experimentation increases the probability of ending up in an inefficient outcome. The more inefficient this outcome, the smaller incentives for experimentation.*

Summing up, any tax cut introduced in the previous regime L is reverted at the onset of a regime H because the dynamic effects of the informational role of policy are irrelevant after a switch, so $\tau_1^* = \bar{\tau}_H^*$ (Proposition 3). According to (R6), the diffusion of θ_t spreads out beliefs μ_t , so switching probability π_t increases with time. When this probability reaches some level, a small tax cut is implemented, turning the switching probability to zero. Tax increments are ruled out from the set of policy options (Proposition 6), so tax cuts are permanent along the regime. When switching probability reaches a higher level, a new tax cut is implemented. Thus, taking into account the diffusion process of the fundamental θ_t , there is a sequence of small and permanent tax cuts. However, the convergence of beliefs in Appendix A implies that there is a floor for these tax cuts. When a new pessimistic regime L finally starts, information about the fundamental is precise again, so tax cuts are reverted such that $\tau_t^* = \bar{\tau}_L^*$, starting a new policy cycle.

5.3 Optimal policy without assuming equilibrium uniqueness

We now study modifications to the optimal policy derived above after relaxing the assumption that tax shifts do not break the equilibrium uniqueness.

5.3.1 Successful big pushes may recover equilibrium multiplicity

To illustrate the only relevant case, let us return to the argument for regimes in Section 3.3. Suppose that when a big push is implemented in a regime L , i.e., $\tau_{t,L} = \inf \{\tau^L\}$, the ex-post observation of labor participation reveals the exact position of the fundamental θ_t . This event triggers an expectation switch, and in turn, provokes a transition to a regime H . In the next period, according to the optimal policy in Proposition 3, taxes should be raised to the benchmark rate, $\tau_{1,H} = \bar{\tau}_H^* \gg \inf \{\tau^L\}$. However, agents in the cluster work in equilibrium only if the big

push is maintained. But with $\bar{\tau}_H^*$, the reward of working is lower than the one with the big push rate. Thus, given the information publicly available about θ_t , agents in the cluster may either work or not work in equilibrium. This indeterminacy breaks the equilibrium uniqueness. Graphically, the after-tax expected participation reward for an arbitrary cutoff strategy $(1 - \tau_t) E_t [m(A(c, \theta_t))]$ twice crosses the 45° line in Figure 4 if $\tau_t = \bar{\tau}_H^*$, but it does only once if $\tau_t = \inf \{\tau^L\}$.

In the above case, the observation of high or low labor participation is uninformative about θ_t . Hence, restrictions (R4) and (R6) still govern the motion of μ_t according to Bayes' rule. This updating rule introduces the standard bias of beliefs in a regime, which flattens the expected participation reward and thus tends to remove this source of multiplicity.

5.3.2 Contingent tax schemes

This subsection proposes a contingent tax scheme that recovers equilibrium uniqueness with high participation in this situation. The equilibrium set $\zeta^*(\theta_t, \tau_t)$ when the fundamental θ_t is observable in (17) implicitly defines the highest possible tax that can eliminate the pessimistic equilibrium (cutoff $c_L^*(\tau_t)$). The next definition introduces notation for this tax rates as a function of θ_t .

Definition 4 *The highest possible tax rate $\tilde{\tau}_L(\theta_t)$ such that $\zeta^*(\theta_t, \tilde{\tau}_L) = \{c_H^*(\tilde{\tau}_L)\}$ is implicitly defined by $c_L^*(\tilde{\tau}_L) = \theta_t$, i.e.*

$$\tilde{\tau}_L(\theta_t) = 1 - \frac{\theta_t}{\varepsilon + (1 - \varepsilon - \varsigma)\beta\theta_t}$$

Notice that $\tilde{\tau}_L$ strictly decreases in θ_t because $(1 - \varepsilon - \varsigma)\beta < 1$. The government could use a design contingent on which equilibrium takes place $\{\tilde{\tau}_L(\theta_t), \bar{\tau}_H^*\}$, such that $\tilde{\tau}_L(\theta_t)$ is levied *if the low participation equilibrium takes place*, and the static optimal rate $\bar{\tau}_H^*$ is levied *if the high participation equilibrium takes place*. By definition, the rate $\tilde{\tau}_L(\theta_t)$ eliminates the low participation equilibrium. Thus, the equilibrium delivered by that tax scheme is unique and efficient since the participation is high and the tax rate attains the optimal provision of the public good. Intuitively, the government can use its commitment technology to ensure that participation is profitable for all agents with $c \leq c_H^*(\bar{\tau}_H^*)$: If the cluster does not participate, $\tilde{\tau}_L(\theta_t)$ is low enough to make them deviate; and if the cluster participates, it is still optimal to do so with taxes $\bar{\tau}_H^*$.

However, the drawback of such policy is that it is feasible only at the onset of an optimistic regime. Because there is no public debt, a tax $\tilde{\tau}_L(\theta_t) < 0$ is not feasible, even with commitment technology.²⁰ Moreover, $\tilde{\tau}_L(\theta_t)$ is negative for $\theta_t > c_L^*(0)$. The information available at the onset of a regime L is that $\theta_{t-1} > c_H^*(\tau_t) - \sigma$, which is higher than $c_L^*(0)$ for any $\tau_t \leq \bar{\tau}_H^*$ by condition (29). Hence, the contingent tax scheme that delivers efficiency is not feasible. This result is strengthened when time elapses in the regime L , since beliefs are biased towards pessimism. However, the

²⁰No public debt is a strong assumption, but the basic point here is that, even if debt is possible, it is not realistic to assume that the government can offer an arbitrary amount of subsidy to everybody in an economy.

information revealed at the onset of a regime H after a big push is $\theta_t < c_H^* (\inf \{\tau^L\}) - \sigma$, in which case such contingent plan is feasible.²¹

The next proposition modifies Proposition 3 to pin down optimal taxes at the onset of regimes taking into account this improvement.

Proposition 7 *Optimal policy at the onset of a regime H is a contingent tax scheme $\left\{ \tilde{\tau}_L \left(w_{(\inf \{\tau^L\})} \right), \bar{\tau}_H^* \right\}$ which removes equilibrium multiplicity if a big push has been successful in the previous period to trigger an expectation switch.*

The rest of our results remain unchanged with respect to our analysis when equilibrium uniqueness is imposed as an assumption. The optimal tax scheme follows the cycle described above according to propositions 4 through 7.

6 Conclusion

This paper studies taxation on labor income in an economy without capital as a vehicle to explore the *informational role* of policy in environments where occasional large fluctuations are driven by expectations switches. The ingredients for such switches to arise are strategic complementarity in payoffs, heterogeneous information about underlying fundamentals, and shocks to fundamentals. Dynamics of these economies are inefficient, since expectations could coordinate in pessimism, despite the fact that optimism is also possible for the same fundamentals. We focus on a case in which the government has no control and no information superiority about fundamentals. Our aim is to shed light on: (1) the channels in which the informational role of policy operates; (2) the policy trade-offs involved; (3) the optimal policy and its contrasting features with standard prescriptions; and (4) the implications of following these standard prescriptions. The primary findings are as follows:

The channels. Distortionary policy affects expectations about what others do and the accumulation of shocks that can trigger expectation switches. These leverages translate into effects of policy on agents' decisions given a specific equilibrium, on the probability of expectation switches, on what agents can learn about fundamentals after observing aggregate outcomes, and on producing non-fundamental volatility if policy breaks equilibrium uniqueness.

Policy trade-offs. An expansionary policy, such as cutting taxes, decreases the risk of falling into pessimism during optimism and increases the chances of an expectation switch during pessimism. However, such a policy often generates unfavorable information about the fundamentals. In a pessimistic regime, if the policy triggers a switch, it feeds expectations that the duration of the upcoming optimism is shorter; and if the policy fails to do so, it is then less likely for a switch

²¹With a similar logic, the government could implement a contingent policy on the onset of regimes L that could force equilibrium multiplicity. However, we abstract from this possibility because, in a more general setting with risk averse agents, it is unlikely that a government-induced non-fundamental volatility could be socially desirable.

to optimism to occur in the near term. In an optimistic regime, an expansionary policy tends to increase the expectations of a switch in the near term. For a contractive policy, such as raising taxes, similar arguments apply but with an opposite sign. Moreover, large tax shifts may recover the equilibrium multiplicity.

Optimal policy. The informational role of policy is a leverage strong enough to motivate a stabilization policy based on the coordination of expectations. After an expectation switch, in either direction, policy should concentrate on its other objectives, rather than on the stabilization. This is because the probability of a new expectation switch is very low in the near term. But with time, when there is a higher probability of a switch, a gradual and permanent expansionary policy is efficient to extend the duration of an optimistic regime, and a severe and transitory expansionary policy is efficient to break pessimism. A contingent policy on which equilibrium takes place is useful for eliminating non-fundamental volatility when large shifts in policy at the onset of optimism deliver equilibrium multiplicity.

Implications of standard prescriptions. Neoclassical theory stresses the desirability of policy that is dynamically neutral, such as time-invariant taxes. However, in economies with self-fulfilling dynamics, such a prescription does not correct the inefficiency resulted from the dynamic coordination failure. New-Keynesian stabilization policy does not correct the inefficiency either. This is because if policy acts immediately after expectation switches, it has little effect on smoothing fluctuations, but produces large costs in momentary welfare. And policy experimentations to extract information about the fundamentals increase the risk of falling into pessimism or staying there for a longer time. Hence, incentives for such policy strategies are smaller in our context than in economies where regimes and switches are assumed exogenous.

Future research should extend our results to build on understanding of the informational role of policy in other macroeconomic and financial environments, where dynamics contain a relevant self-fulfilling component and are affected by occasional expectation switches.

Appendix

A. Proofs and derivations

Convergence of beliefs conditional on staying in the low activity equilibrium without taxes is:

$$\lim_{t \rightarrow \infty} \mu_t = \begin{cases} 0 & \text{for } w_i < c_L^* \\ 2 \tan\left(\frac{1}{2}r\right) \sin(r(i-H)) & \text{otherwise} \end{cases}$$

where $r = \frac{\pi}{2(K-L)+1}$ and L implicitly defined by $c_L^* \in [w_L, w_{L+1}]$.

(from Chamley [10]) Given the process on μ_t , convergence implies

$$\mu_L = p\mu_{L+1} \iff \mu_{L+1} = \frac{\mu_L}{p} = a$$

$$\mu_{L+i} = \frac{p\mu_{L+i-1} + (1-2p)\mu_{L+i} + p\mu_{L+i+1}}{1-\mu_L}$$

$$\mu_{L+i+1} - (2-a)\mu_{L+i} + \mu_{L+i-1} = 0 \text{ for } i \in [L+2, K-2]$$

$$\mu_{K-1} = \frac{\mu_K}{1-a}$$

or equivalently, $\mu_{K-1} = \mu_K$. The solution for the polemonium is complex, satisfying $\mu_{L+i} = C \sin(r(i + \varkappa))$. The special cases allow to find these parameters,

$$\begin{aligned} C \sin(\varkappa) &= 0 \implies \varkappa = 0; \quad \mu_{L+1} = a \implies C \sin(r) = a; \\ \mu_{L+2} &= (2-a)\mu_{L+1} \implies \sin(2r) = (2-a)\sin(r); \\ \mu_{K-1} &= \mu_K \implies \sin(r(K-L)) = \sin(r(K+1-L)) \end{aligned}$$

Taking into account that $\mu_{L+i} \geq 0$ and the last condition,

$$\begin{aligned} \sin(r(K-L)) &= \cos\left(\frac{\pi}{2} - r(K-L)\right) = \cos\left(r(K+1-L) - \frac{\pi}{2}\right) = \sin(r(K+1-L)) \\ \frac{\pi}{2} - r(K-L) &= r(K+1-L) - \frac{\pi}{2} \iff r = \frac{\pi}{2(K-L)+1} \end{aligned}$$

from the rest of the conditions,

$$a = 2 - \frac{\sin(2r)}{\sin(r)} = 2 - \frac{2\sin(r)\cos(r)}{\sin(r)} = 2(1 - \cos(r))$$

$$C = \frac{a}{\sin(r)} = \frac{2(1 - \cos(r))}{\sin(r)} = 2 \left(\frac{1 - \cos^2(r/2) + \sin^2(r/2)}{2\cos(r/2)\sin(r/2)} \right) = 2 \tan(r/2)$$

Proof of Proposition 4. An increase in taxes $\tilde{\tau} \in \tau^L$ s.t. $\tilde{\tau} > \bar{\tau}_L^*$ decreases one-period welfare and turns the switching probability $\pi_t(c_L^*(\tilde{\tau})) = 0$ because $w_L(\tilde{\tau}) < w_L(\bar{\tau}_L^*)$, so there are no

incentives to increase taxes above $\bar{\tau}_L^*$. If taxes are cut, i.e. $\tilde{\tau} \in \tau^L$ s.t. $\tilde{\tau} < \bar{\tau}_L^*$, then the effects on expected welfare are (compared to the case when $\bar{\tau}_L^*$ is levied): (i) decreases the one-period welfare; (ii) increases $\pi_t(c_L^*(\tilde{\tau}))$; (iii) if there is no switch, expectations of a future switches decreases (effects on μ_{t+1}), with a negative effect on welfare; and (iv) if there is a switch, the expected duration of the upcoming optimistic regime becomes shorter (by the effect on μ_1), also with a negative effect on welfare. Thus, the only incentive to implement a tax cut is to increase the switching probability. However, as the feasible tax rate is bounded below, there is an upper limit of the reachable switching probability $\bar{\pi}$, critically depending on the shape of the lower tail of the beliefs μ . If $\bar{\pi}$ is too low, a tax cut may generate too little gain compared to the costs associated with effects (i), (iii) and (iv). Hence, there exists some threshold on $\bar{\pi}$ that triggers a tax cut.

Once a tax cut is triggered, we argue that it will reach a corner solution, $\tau_t^* = \inf \{\tau^L\}$, which we label as a "big push". There are two reasons for this. First, the marginal effect of tax cuts on switching probability $\pi_t(c_L^*(\tau_t))$ increases in the size of such tax cuts. It is because the evolution of beliefs assign higher probability to higher positions at the lower tail (for instance, as in Figure 3B.) And more severe tax cut moves the threshold $w_{L(\tau_t)}$ up by more positions.

Second, such increasing return in raising switching probability is weighted by the difference in welfare between the two regimes, $V(\mu_1) - W(\mu_{t+1})$ in the efficiency condition (27). However, other welfare effects of tax cuts are either limited by the change of one-period welfare (effect i), or they are quite smooth because the effect (iii) and (iv) depends on future optimal policies (a discrete version of the envelope theorem.) Therefore, if the difference in participations between regimes is large enough (denoted by the condition $\tilde{\kappa}$), the substantial welfare gap across regimes implies a dominant welfare effect of increasing switching probability by tax cuts. It then rationalizes the corner solution $\tau_t^* = \inf \{\tau^L\}$ in tax cuts. ■

Proof of Proposition 6. Any tax rate $\tilde{\tau} \in \tau^H$ s.t. $\tilde{\tau} < \tau_{t-1}^*$ turns the switching probability $\pi_t(c_H^*(\tilde{\tau})) = 0$ because $w_H(\tilde{\tau}) < w_H(\tau_{t-1}^*)$. And if $\tau_{t-1}^* \leq \bar{\tau}_H^*$, a smaller $\tilde{\tau}$ delivers lower one-period welfare. Thus, if a tax cut is optimal to implement, the tax rate is the nearest lower position to τ_{t-1}^* , i.e. $\tau_t^* = \left\{ \tau_k^H : w_H(\tau_{t-1}^*) - w_H(\tau_k^H) = v \right\}$.

We want to rule out any possible tax increment as an optimal policy in the optimistic regime.

First, let's look at an increase in taxes to the nearest higher position. Such policy produces the following effects: (i) increases one-period welfare (if $\tau_{t-1}^* < \bar{\tau}_H^*$); (ii) increases transition probability, with negative effect on welfare weighted by the difference in value between the two regimes, $V(\mu_{t+1}) - W(\mu_1)$, as shown from the efficiency condition (28); (iii) if the regime continues, future switching probability decreases (by the effect on μ_{t+1}), with a positive impact on welfare; and (iv) if a switch occurs, the expected duration of the upcoming "pessimistic" regime becomes shorter (by an effect on μ_1), with a positive effect on welfare.

If a tax cut has been implemented in the past, a tax increase to the nearest higher position cannot be the dominant strategy in the near term. Such a policy implies a switching probability

$\pi_t = \mu_{H(\bar{\tau}),t} + p\mu_{H(\tau_{t-1}^*),t}$. However, a previous tax cut at some period $t - j$ (from a tax rate $\bar{\tau}$) was implemented to avoid a possible switch to the pessimistic regime with probability $p\mu_{H(\bar{\tau}),t-j}$. That means the loss in welfare by a possible switch (effect ii) dominated other benefits (effects i, iii and iv) by having the tax rate equal or higher than $\bar{\tau}$ at $t - j$. Such net welfare effect of cost-benefit trade-off holds from $t - j$ to t for two reasons. First, the cost of increasing the tax rate to the nearest higher position is larger because $p\mu_{H(\bar{\tau}),t-j} < p\mu_{H(\tau_{t-1}^*),t} < \pi_t$. It is because the evolution of beliefs in one regime assigns higher and higher probabilities with time t to the same positions at the upper tail of the distribution (Figure 3B), hence $\mu_{H(\bar{\tau}),t-j} < \mu_{H(\bar{\tau}),t}$. Moreover, beliefs assign higher probabilities to lower locations at the upper tail of its distribution so that $\mu_{H(\bar{\tau}),t} < \mu_{H(\tau_{t-1}^*),t}$ as $\tau_{t-1}^* < \bar{\tau}$ due to previous tax cuts. Second, the benefit side does not change much with time, only with a bigger gain by tax increments in the one-period welfare because the tax rate is now lower than before. However, such increased gain is limited in the one-period welfare and is hard to compete with the cost in switching probability that is weighted by the welfare gap across two regimes. In sum, if there has been a tax cut, it won't be desirable to have a small tax increase.

If no tax cut has been implemented in the past, then a tax increase also decreases one period welfare (it is a deviation from $\bar{\tau}_H^*$). Thus, this possibility is ruled out.

Then we turn to the case with a large tax increase. A higher tax increment increases benefits (iii) and (iv) from learning. However, these two benefits increase at a decreasing rate with the size of the tax increment because they depends on future optimal policies (a discrete version of the envelop theorem). In addition, the effect (iv) must be smooth because the government cannot perfectly infer at which location the fundamental will be revealed, so the computation of expectations on this term introduces additional decreasing returns of such a policy. In the meantime, the negative effect in switching probability increases with the size of the tax increment, so a policy of large tax raising is ruled out.

Thus, the only policy options is to continue with the tax levied on the period before, or reduce taxes to the nearest lower position in the set of admissible policies τ^H . ■

B. A Numerical example

This appendix numerically solves the optimal time path of taxes for the problem in (24) and (25) subject to (R1) through (R6), which assumes equilibrium uniqueness and is studied in Sections 5.2 and 5.3. We compare the efficiency of the resulting tax scheme with a time-invariant policy focused only on short run objectives (levying $\bar{\tau}_L^*$ or $\bar{\tau}_H^*$, the static optimum tax rate in the respective regime). This exercise allows us to evaluate the dynamic effects of the informational role of policy in a simpler environment than in the full model, where further improvement taking into account equilibrium multiplicity considers the introduction of a contingent design along the lines of Proposition 7.

To compute welfare, we need to specify a functional form for the public good utility,

$$\phi(g) = \chi_s \frac{g^{1-\rho}}{1-\rho}, \text{ where } s = L, H \quad (30)$$

The parameter χ_j helps to simplify our exercise, since it allows to calibrate the static optimum tax rates to be equal in both regimes, i.e. $\bar{\tau}_L^* = \bar{\tau}_H^* = \bar{\tau}^*$.

B.1 The candidate plan and its benchmarks

Although our policy problem is well defined as a Markov chain with two regimes, its solution using numerical methods involves the challenge that the transition matrix is endogenous to policy, so the state is a vector μ representing beliefs, which has a large dimension. Hence, in our solution, we need to deal with "the curse of dimensionality", which in our case cannot be rounded by using a few parameters to approximate the function describing beliefs. This is because μ_t starts from a degenerated *pdf* with a single point concentrating the whole mass right after a transition ($t = 0$), and converges to a cosine function if the regime is maintained ($t \rightarrow \infty$, for example, see Figure 3B)

Therefore, to solve this problem, we must depart from popular algorithms, such as Krusell and Smith [20]. We tackle this difficulty by using the specific properties of our problem, as stated in Propositions 2 through 6. Thus, we translate our dynamic programming problem with a continuous control variable into an optimal waiting problem with two options each period for each regime.

A detailed description of the solution algorithm is displayed in the Appendix C, but the basic intuition is the following. Proposition 3 pins down the optimal tax at $t = 1$, $\tau_{1,s}^* = \bar{\tau}^*$ for $s = L, H$ ($\bar{\tau}_L^* = \bar{\tau}_H^*$ by calibration). For $t > 1$, we use Propositions 4 and 6 to sequentially evaluate the two surviving policy options in regime L (the static optimum $\bar{\tau}_L^*$, or the big push, $\inf \{\tau^L\}$), and in regime H (constant taxes, $\tau_t = \tau_{t-1,H}^*$, or cutting them in one step). To compute the value of these options, we assume a sequence of future taxes $\left\{ \tau_{i,s}^* \right\}_{i=t+1}^T$ for $s = L, H$ and a terminal period T large enough so the diffusion of θ_t (R6) becomes irrelevant for μ_t . Later we obtain the value functions $V(\cdot)$ and $W(\cdot)$ for all possible initial conditions μ_1 , and use them to iterate upon exact convergence of $\left\{ \tau_{i,s}^* \right\}_{i=1}^T$ and approximate convergence for $V(\cdot)$ and $W(\cdot)$.

B.2 Results

Table 1 summarizes the calibrated parameters, which satisfies condition (29). Figure 5 shows one period welfare as a function of taxation, with $\bar{\tau}^* = 30\%$. To evaluate the incidence of strategic complementarity in individual decisions on the marginal effect of taxes on one period welfare, we shut down the externality in the participation reward function, so $m(A_t) = \bar{m}$, which is calibrated to match $m(A_L^*)$ and $m(A_H^*)$ with $\bar{\tau}^* = 30\%$. The optimal tax rate in this case is much higher: 40% for regime L and 70% for regime H . Proposition 1 anticipates this result, since it states that

the complementarity amplifies the (negative) marginal effect of taxes on real activity and that this effect is stronger when participation is high.

Participation in the calibrated economy with taxes $\bar{\tau}^*$ is $A_L = 2.5\%$ ($c_L^* = 8.5\%$) and $A_H = 87\%$ ($c_H^* = 0.56$), since the cluster is 70% of the total population. As result, the reward of participation in the optimistic regime is about four times higher than in the depressed one.

We assume that $\theta_t \in [0, .8]$ (since $\sigma = .2$), and its partition $\{w_k\}_{k=1}^K$ has $K = 41$ steps, each with length $v = .02$. Under these conditions, $c_L^*(\tau_t) \in [0, .12]$, so $\tau^L = \{6\%, 19\%, 30\%\}$. We discard taxes higher than $\bar{\tau}^*$ since those taxes are always suboptimal in the regime L . For regime H , $c_H^*(\tau_t) \in [.2, .96]$, so τ^H has 30 elements with tax rates between 0% to 60%. The maximum tax is 60%, since this is the minimum rate that triggers a transition for sure, i.e., that the cluster does not participate in equilibrium for any θ_t .

Before we start presenting our numerical results, we need to remark that time t in our model has two twisted interpretations: it represents time elapsed from the last transition; and its frequency is determined by how often the economy is hit by small shocks, which could have weekly, daily or even higher frequency depending on the application. Therefore, our results to be presented below have no direct interpretation in available time series for most macroeconomic data.

B.2.1 Taxes

Figure 6 illustrates the sequence of optimal taxes in each regime for two examples of initial information: Figure 6A assumes that taxes during the last transition was $\bar{\tau}^*$; and Figure 6B assumes those taxes at the minimum rate in equilibrium. The left figures show optimal taxes for the regime L (pessimistic phase) and the right figures for the regime H (optimistic phase).

In all cases, optimal taxes sharply depart from our benchmark of a time-invariant tax policy $\bar{\tau}^*$. As anticipated by Proposition 3, optimal taxes on the onset of both regimes coincide with the benchmark $\bar{\tau}^*$, since policy has little power on expectations when information about the fundamental θ_t is precise, which is the case in these situations. For later periods, transition probability increases as information about the fundamental θ_t gets diffused, so when this probability reaches some threshold, the authority deviates from the benchmark policy in the ways anticipated by Propositions 4 to 6, with transitory big pushes in a regime L , and permanent small tax cuts in a regime H . However, the authority waits longer time for implementing an active policy during regime L than during regime H because it has to give up more of its short run objectives in a regime L .

Taxes in a regime L therefore show an oscillating path (left panels of Figures 6A and 6B), where a big push implies decreasing taxes to 6%, and if that attempt to break pessimism fails, the information revealed reverts such policy, so the government levies again $\bar{\tau}^* = 30\%$. In the regime H (right panels of Figures 6A and 6B), since tax cuts are permanent as long as this regime survives, the sequence of optimal taxes decreases slowly in steps, each time with a longer waiting time because each tax cut implies higher cost in terms of one period welfare and in learning. In

this numerical example, costs of a new tax cut are higher than its benefits when taxes reach 13%, so a time-invariant tax policy is implemented after that point.

In a regime L , but with higher taxes during the last transition (comparison between left panels of Figure 6A and 6B), the probability of a new transition increases faster, so the waiting time for the first push is shorter, as also does so the waiting period between pushes. For the regime H (right panels of Figure 6A and 6B), the government takes longer to deviate from $\bar{\tau}^*$ when taxes were higher during the last transition because, in this case, that information implies that transition probability increases more slowly with time.

B.2.2 Transition probability and welfare

Figure 7 and Figure 8 summarize the equilibrium dynamics with optimal taxes (solid line) and with the benchmark time-invariant policy $\bar{\tau}^*$ (stars line). Figure 7 shows transition probability and Figure 8 shows cumulative transition probability (also called "hazard"), with panels A and B respectively for the optimal policy sequences of Figure 6A and 6B, which have different initial information.

In a regime L the spikes in Figure 7 (left graphs) capture the high transition probability attained by each big push and the respective depression in expectations if those pushes fail. The probability reached during the push increases with time as beliefs μ_t about θ_t become more diffused. Figure 8 shows that the hazard of a transition increases faster with the active policy, reaching 35% after 300 periods instead of 25% in the panel A, despite the depressed expectations after unsuccessful pushes partially offset the benefit of such policy. The comparison between the left graphs of panels A and B of Figures 7 and 8 shows that the higher transition probability obtained because taxes were higher during the last transition implies that policy in the current regime L is also more efficient (a gain of 10% instead of 2% after 300 periods).

In the regime H , the active plan has significantly lower transition probability with respect to the benchmark time-invariant policy $\bar{\tau}^*$ (comparison between solid and star lines in right graphs of Figure 8): around 25% lower after 300 periods. Figure 7 shows why: each tax cut turns the transition probability to zero, even though it increases quickly after that.

In regard to welfare, the evolution of the expected transition probability shapes its time path. In the regime L , welfare is low, but the gain obtained by the active plan is very relevant, with an impact around 30% in the calibrated example. When Figure 9A is compared to Figure 9B, we observe that welfare is higher when taxes on the period of the last transition are also higher, suggesting that the effect of taxes on learning is far from being irrelevant. For a regime H , the right graphs of Figure 9 show that welfare decreases as transition probability increases, but the active plan delivers approximately 9% higher welfare than the benchmark policy that ignores dynamic effects of its informational role. Since the effect of taxes on learning operates in this regime H in the opposite way as in regime L , higher taxes during the last transition implies lower welfare, since,

given the initial information revealed, transition probability increases more quickly.

B.3 Optimality of our candidate plans

As with any numerical methods, we can only approximate a fully optimal policy. This is specially important in our case, where policy decisions are restricted to be only two and where we need to assume an optimal future sequence of taxes for each history and decision in each period. However, what we can show is that our candidate tax plans are equilibrium strategies, in the sense that the government would not deviate to other policy if it can choose a third policy option given private agents' equilibrium strategies.

The difference in welfare between regimes, the endogeneity of the transition probability, and the properties stated in Proposition 6 ensures that the two options for policy and our iteration algorithm really summarize the policy problem in the regime H . However, our numerical strategy for the regime L is more arguable, since the two surviving options depend on the fact that the optimal big push is really a corner solution in the set τ^L of policy options (Proposition 4).

To check this condition, Figure 10 shows the difference in welfare in the regime L between our candidate solution and an alternative policy option. This alternative policy includes a deviation to an "intermediate push," so the transition threshold $w_L(\tau_t)$ is taken to the middle position (recall that in this calibration τ^L has only three elements), i.e. that $\tau_{t,L} = 19\%$. Figure 10A shows the result for the candidate tax plan in Figure 6A, and Figure 10B for the tax plan in Figure 6B. In both cases, the value of the candidate plan is always higher than the value of a "moderate push" plan, for all periods in the time window used for this numerical exercise. Therefore, we conclude that our numerical strategy delivers a candidate tax plan, which is equilibrium.

C. Solution Algorithm

Step 1 Start from a guessed $V_1^{L(i)}(\theta_0^L)$ vector, which is the low phase value at the period right after a switch conditional on different starting position of cluster. i is the index for iteration.

Step 2 For each starting position of cluster θ_0^H , compute the value of high phase at the period right after a switch: $V_1^{H(i)}(\theta_0^H)$

- (a) The initial likelihood function F_0 is set to have a mass 1 on the position of θ_0^H .
- (b) Compute forward the optimal tax choice from the set τ^H defined in proposition 13. Sequentially, choose the maximum value from the following 3 optional strategies available at each period:

1st iteration

- i. lower tax rate by one step and keep time-invariant tax rate at that level forever (assuming convergence after 700 periods)

- ii. keep the current tax rate the same as last period but lower it by one step next period and keep constant there forever
 - iii. set the tax rate at the highest possible level to push the economy into low phase with certainty
- (c) We use the values at current period implied by these 3 options to select the optimal tax rate: $\tau_t^* = \begin{cases} \max \{ \tau < \tau_{t-1}^* | \tau \in \tau^H \} & V_t^i \text{ is biggest} \\ \tau_{t-1}^* & V_t^{ii} \text{ is biggest} \\ \tau_0 & V_t^{iii} \text{ is biggest} \end{cases}$

The value is computed using (12) and as described by steps *g* and *h*. Note that for different tax rate, the inferred position of the cluster once a switch occurs varies accordingly, which implies a different value of the low activity phase if a switch occurs at the tax rate: $V_1^{L(i)}(\theta_0^L(\tau_t))$

2nd iteration and on upon convergence

Keeping the policy time path obtained from the previous iteration $\Gamma^{(k-1)} = \left\{ \tau_t^{(k-1)} \right\}_{t=1}^T$, evaluate the following 3 optional strategies:

- i. lower tax rate by one step, assuming for the future path the one suggested by the $\Gamma^{(k-1)}$ from the time where the tax rate coincides with the current option (assuming convergence after 700 periods)
 - ii. keep the current tax rate one period and apply the same tax path as option *i* from next period and on
 - iii. set the tax rate at the highest possible level to push the economy into low phase with certainty
- (d) Use the values implied by these 3 options in identical manner as step *c*
- (e) Iterate upon exact convergence in the policy sequence
- (f) Get the corresponding components implied by the optimal tax rate, which include: probability of switch $\pi(\theta_0^H, \tau_t^*)$, momentary payoffs $A^H(\tau_t^*)$, the likelihood function $F_{t+1}^H(\tilde{\omega}; \theta_0^H, \tau_t^*)$ and the value of low phase if a switch occurs $V_1^{L(i)}(\theta_0^L(\tau_t^*))$.
- (g) After an arbitrarily long period T (700 in our numerical exercise), we assume V_T^H will prevail in the rest of time. Therefore,

$$V_T^H = \frac{A^H(\tau_T^*) + \delta \pi(\theta_0^H, \tau_T^*) V_1^{L(i)}(\theta_0^L(\tau_T^*))}{1 - \delta(1 - \pi(\theta_0^H, \tau_T^*))}$$

- (h) Using V_T^H and the stream of $\{\pi_t\}_{t=1}^T, \{A_t\}_{t=1}^T, \{F_t\}_{t=1}^T, \left\{ V_{1t}^{L(i)} \right\}_{t=1}^T$ computed at step *c*, we go backwards until period 1 to get the high phase value associated with the initial position of cluster $V_1^{H(i)}(\theta_0^H)$

Step 3 Having got the vector $V_1^{H(i)}(\theta_0^H)$, we adopt a similar strategy as in Step 2 to compute an updated vector $V_1^{L(i+1)}(\theta_0^L)$.

- (a) For each starting position of cluster θ_0^L , the initial likelihood function F_0 is set to have a mass 1 on the position of θ_0^L .
- (b) Compute forward the optimal tax choice from the set τ^L defined in proposition 13. This time, in low phase, only 2 optional strategies is available at each period:

1st iteration

- i. set the lowest tax in the set τ^L in the hope to trigger a switch. In next period, set the tax rate back to the static optimal level forever
 - ii. currently keep the static optimal tax rate but set the lowest tax in next period. After that trial, reset the tax rate at static optimal level forever
- (c) We also use values at current period implied by these 3 options to select the optimal tax rate:

$$\tau_t^* = \begin{cases} \tau_1 \in \tau^L & V_t^i \text{ is biggest} \\ \tau_s \in \tau^L & V_t^{ii} \text{ or } V_t^{iii} \text{ is biggest} \end{cases}$$

The value is computed using (12) and as described by steps *g* and *h*. Note that for different tax rate, the inferred position of cluster once a switch occurs need to vary accordingly, which implies a different value of high phase if a switch occurs at *the* tax rate: $V_1^{H(i)}(\theta_0^H(\tau_t))$

2nd iteration and on upon convergence

Keeping the policy time path obtained from the previous iteration $\Gamma^{(k-1)} = \left\{ \tau_t^{(k-1)} \right\}_{t=1}^T$, evaluate the following 2 optional strategies:

- i. set the lowest tax in the set τ^L in the current time period, assuming for the future path the one suggested by the $\Gamma^{(k-1)}$ from the time where the tax rate coincides with the current option (assuming convergence after 700 periods)
 - ii. currently keep the static optimal tax rate but set the lowest tax in next period and apply the same tax path as option *i* from next period and on
- (d) Use the values implied by these 2 options in identical manner as step *c*
- (e) Iterate upon exact convergence in the policy sequence
- (f) Get the corresponding components implied by the optimal tax rate, which include: probability of switch $\pi(\theta_0^L, \tau_t^*)$, momentary payoffs $A^L(\tau_t^*)$, the likelihood function $F_{t+1}^L(\tilde{\omega}; \theta_0^L, \tau_t^*)$ and the value of low phase if a switch occurs $V_1^{H(i)}(\theta_0^H(\tau_t^*))$.
- (g) After an arbitrarily long period T (700 in our numerical exercise), we assume V_T^L will prevail in the rest of time. Therefore,

$$V_T^L = \frac{A^L(\tau_T^*) + \delta \pi(\theta_0^L, \tau_T^*) V_1^{H(i)}(\theta_0^H(\tau_T^*))}{1 - \delta(1 - \pi(\theta_0^L, \tau_T^*))}$$

- (h) Using V_T^L and the stream of $\{\pi_t\}_{t=1}^T$, $\{A_t\}_{t=1}^T$, $\{F_t\}_{t=1}^T$, $\left\{ V_{1t}^{H(i)} \right\}_{t=1}^T$ computed at step c),

we go backwards until period 1 to get the high phase value associated with the initial position of cluster $V_1^{L(i+1)}(\theta_0^L)$

Step 4 Check convergence in vector $V_1^L(\theta_0^L)$. If the vector converges, so does the vector $V_1^H(\theta_0^H)$. And we collect the streams of optimal tax rates $\{\tau_t^{H*}, \tau_t^{L*}\}_{t=1}^T$, the probabilities of switch $\{\pi_t^{H*}, \pi_t^{L*}\}_{t=1}^T$ and the values $\{V_t^H, V_t^L\}_{t=1}^T$ for both phases and for each possible initial position of cluster. If it does not, use $V_1^{L(i+1)}(\theta_0^L)$ obtained in step 3 and iterate again from step 2.

Step 5 Compute the cumulated probability of switch using $\{\pi_t\}_{t=1}^T$ by



$$\pi_t^c \equiv \text{prob}(\text{switch at } t \mid \text{no switch occurs until } t-1) = \pi_t \prod_{s=1}^{t-1} (1 - \pi_s)$$

$$\pi_t^c \equiv \text{prob}(\text{switch occurs in the past } t \text{ periods}) = \sum_{s=1}^t \pi_s^c$$

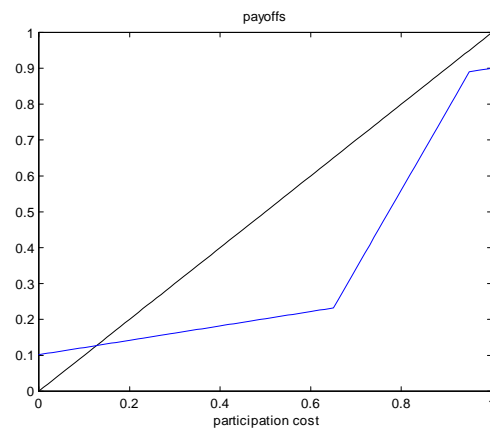

D. Tables and figures

Parameter:	Denoting:	Value
ε	Lower bound on participation payoff	.1
ς	Upper bound on participation payoff	.1
α	Parameter of <i>pdf</i> inside the cluster	3.5
β	Parameter of <i>pdf</i> in the complete support	.3
σ	Length of cluster	.2
p	Probability on the random walk for θ	$\frac{1}{4}$
v	Length of step in partition of θ	.02
$\underline{\tau}$	Minimum tax rate allowed	0
$\bar{\tau}$	Maximum tax rate allowed	.6



Table 2 – Parameters for simulations



A - Two equilibria



B - One equilibrium with low participation



C - One equilibrium with high participation

Figure 1 – Participation payoff and equilibria for alternative θ





Figure 2 – Expected participation payoff at $t + 1$ given observability of θ at t



A - distribution of θ

B - beliefs about θ

Figure 3 – Diffusion process







Figure 4 – Equilibria conditional on $\theta_t = w_L$



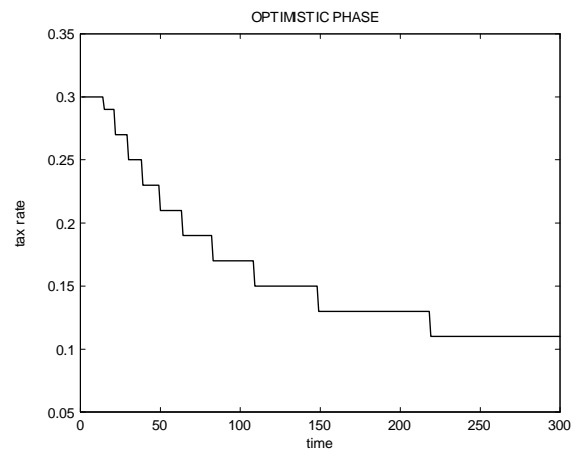
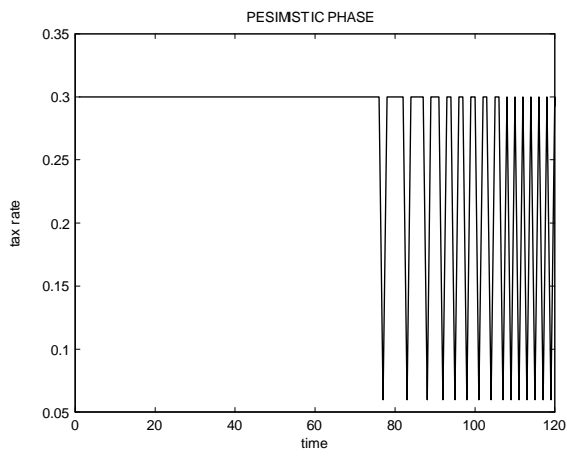
A – "Pessimistic" regime

B – "Optimistic" regime

Figure 5 – Static optimum tax rate


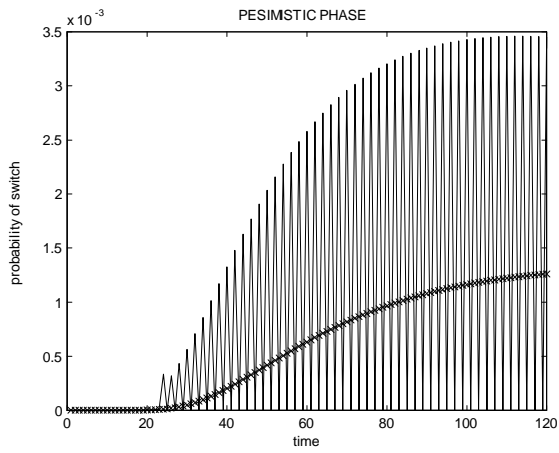


A – When switch triggered at static optimum rate in the alternative regime





B – When switch triggered at minimum rate in the alternative regime

Figure 6 – Policy at the equilibrium path

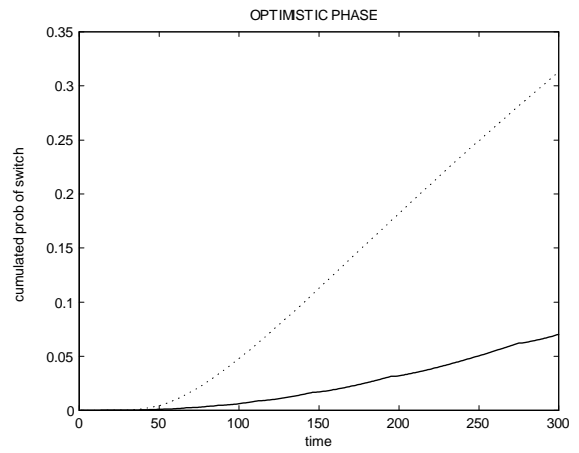



A – When switch triggered at static optimum rate in the alternative regime





B – When switch triggered at minimum rate in the alternative regime

Figure 7 – Switching probability at the equilibrium path for policy (solid) and time invariant policy at static optimum rate (dashed)





A – When switch triggered at static optimum rate in the alternative regime





B – When switch triggered at minimum rate in the alternative regime

Figure 8 – Cumulative switching probability along the equilibrium path for policy (solid) and time invariant policy at static optimum rate (dashed)

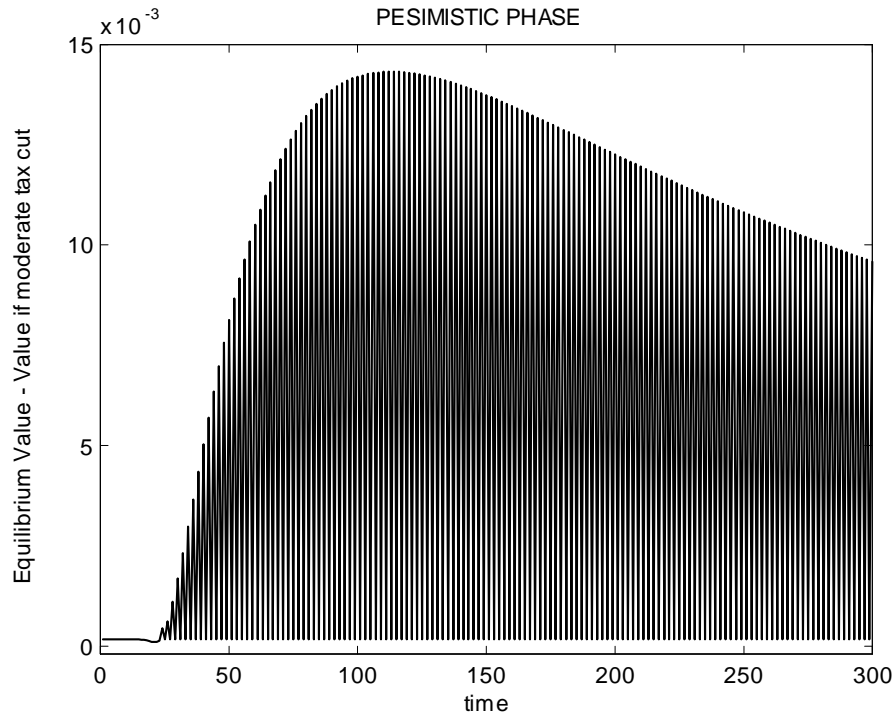


A – When switch triggered at static optimum rate in the alternative regime

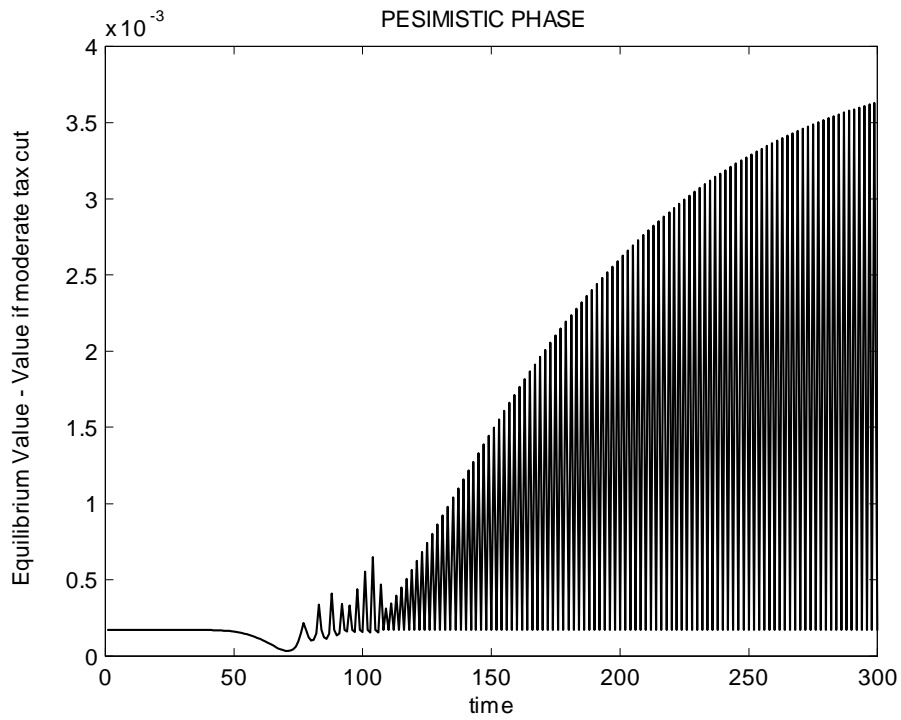


B – When switch triggered at minimum rate in the alternative regime

Figure 9 – Welfare along the equilibrium path for policy (solid) and time invariant policy at static optimum rate (dashed)



A – When switch triggered at static optimum rate in the alternative regime



B – When switch triggered at minimum rate in the alternative regime

Figure 10 – Evaluation of deviation from "big push" to "moderate push" in pessimistic regime

References

- [1] Angeletos, Georges-Marios, Christian Hellwig and Alessandro Pavan. 2006. "Signaling in a Global Game: Coordination and Policy Traps." *Journal of Political Economy* 114(3): 452-84.
- [2] Angeletos, Georges-Marios, Christian Hellwig and Alessandro Pavan. 2007. "Dynamic Global Games of Regime Change: Learning, Multiplicity and Timing of Attacks." *Econometrica* 75(3): 711-56.
- [3] Angeletos, Georges-Marios and Alessandro Pavan. 2007. "Policy with Dispersed Information." NBER Working Papers Series, No.13590.
- [4] Angeletos, Georges-Marios and Ivan Werning. 2006. "Crises and Prices: Information Aggregation, Multiplicity, and Volatility." *American Economic Review* 96(5): 1720-36.
- [5] Azariadis, Costas. 1981. "Self Fulfilling Profecies." *Journal of Economic Theory* 25(3): 380-396.
- [6] Blanchard, Olivier Jean and Nobuhiro Kiyotaki. 1987. "Monopolistic Competition and the Effect of Aggregate Demand." *American Economic Review* 77(4): 647-66.
- [7] Caballero, Ricardo, Emmanuel Fahri, and Mohamad Hammour. 2006. "Speculative Growth: Hints for the U.S. Economy." *American Economic Review* 96(4): 1159-92.
- [8] Carlsson, Hans and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica* 61(5): 989-1018.
- [9] Chamley, Christophe. 1986. "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives." *Econometrica* 54(3): 607-22.
- [10] Chamley, Christophe. 1999. "Coordinating Regime Switches." *Quarterly Journal of Economics* 114(3): 869-905.
- [11] Cooper, Russell and Andrew John. 1988. "Coordinating Coordination Failures in Keynesian Models." *Quarterly Journal of Economics* 103(3): 441-63.
- [12] Cooper, Russell. 1998. "Coordination Games: Complementarities and Macroeconomics." Cambridge, UK: Cambridge University Press.
- [13] Diamond, Peter. 1982. "Aggregate Demand Management in Search Equilibrium." *Journal of Political Economy* 90(5): 881-94.
- [14] Ennis, Huberto and Todd Keister. 2005. "Optimal Fiscal Policy under Multiple Equilibria." *Journal of Monetary Economics* 52(8): 1359-1377.
- [15] Frankel, David and Ady Pauzner. 2000. "Resolving Indeterminacy in Dynamic Settings: The Role of Shocks". *Quarterly Journal of Economics* 115(1): 285-304.

- [16] Hamilton, James. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business-Cycle." *Econometrica* 57(2): 357-84.
- [17] Hamilton, James and Marcelle Chauvet. 2006. "Dating Business Cycle Turning Points." In: *Nonlinear Analysis of Business Cycles*, edited by Costas Milas, Philip Rothman, and Dick van Dijk. North Holland: Elsevier.
- [18] Kiyotaki, Nobuhuro and John Moore. 1997. "Credit Cycles," *Journal of Political Economy*, 105(2): 211-48.
- [19] Howitt, Peter and R. Preston McAfee. 1992. "Animal Spirits." *American Economic Review* 82 (3): 493-507
- [20] Krusell, Per and Anthony Smith. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106 (5): 867-96.
- [21] Marcet, Albert and Ramon Marimon. 1994. "Recursive Contracts." Mimeo, Pompeu Frabra.
- [22] Morris, Stephen and Hyun Song Shin. 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88 (3): 587-97.
- [23] Murphy, Kevin M., Andrei Shleifer, and Robert Vishny. 1989. "Industrialization and the Big Push." *Journal of Political Economy* 97 (5): 1003-26.
- [24] Obstfeld, Maurice. 1986. "Rational and Self-Fulfilling Balance-of-Payments Crises." *American Economic Review* 76(1): 72-81.
- [25] Svensson, Lars and Noah Williams. 2007. "Bayesian and Adaptive Optimal Policy under Model Uncertainty." NBER Working Papers Series, No.13414.
- [26] Woodford, Michael. 1989. "Expectations, Financial Constraints, and Aggregate Instability." In: edited by M. Kohn and S. C. Tsiang, *Finance Constraints, Expectations, and Macroeconomics*, New York, NY: Oxford University Press.