A Pendulum Model of Economic Growth *

Zhen Huo † Xilin Yuan ‡ Xiaodong Zhu §

May, 2026

Abstract

This paper develops a "pendulum model" of economic growth to explain the recurrent oscillations between state control and market liberalization observed in major economies. We posit that policymakers face a fundamental trade-off between promoting growth and maintaining political or systemic control over the accumulation of private wealth or capital. Using a dynamic framework where a planner balances these competing objectives, we show that high growth targets endogenously generate cycles: periods of pro-market expansion are necessary to meet growth targets but inevitably increase the share of private resources, triggering a retrenchment phase to restore control. We provide empirical evidence supporting this mechanism, focusing primarily on China's post-reform history, where policy swings between "pro-private" and "pro-state" regimes align closely with the model's predictions. To demonstrate the broader applicability of the trade-off between efficiency and stability, we present supplementary evidence from the United States, illustrating similar cyclical patterns in financial regulation and credit cycles. Our findings suggest that the volatility of economic policy is not merely a response to exogenous shocks but an inherent feature of governance under binding growth constraints.

Keywords: growth targets, endogenous cycles, financial frictions

∗We are grateful for comments from Munseob Lee, Xuewen Liu, and participants at the HKU Brown Bag Seminar, the 2025 CUHK-UCSD Joint Lab Conference, and the 4th China Workshop on Growth and Development for useful feedback.

†Yale University, Department of Economics, zhen.huo@yale.edu

‡Yale University, Department of Economics, xilin.yuan@yale.edu

§The University of Hong Kong, xdzhu@hku.hk

1. Introduction

The trajectory of modern economic development is rarely linear. Instead, it is often characterized by distinct phases of acceleration and deceleration, marked by alternating regimes of state intervention and market liberalization. In China, for instance, the post-1978 reform era has witnessed recurring "pendulum" swings: periods of vigorous non-state sector expansion in the 1980s and product and labor market liberalization between 2001 and 2007 were followed by phases of retrenchment and state consolidation, such as the "governance and rectification" campaign of the early 1990s and the recent reassertion of state control over the platform economy. Similarly, in the United States, financial history exhibits cyclical patterns of deregulation followed by regulatory tightening, often synchronized with credit expansions and financial crises. These observations raise a central question: What drives these systematic oscillations in economic policy, and how do they interact with macroeconomic performance?

In this paper, we propose a theoretical framework—the "pendulum model"—to explain these dynamics. We argue that policy oscillations are endogenous outcomes of a planner's attempt to reconcile two conflicting objectives: meeting explicit growth mandates and maintaining control over the distribution of economic resources. In our model, the planner faces a hard constraint on the share of private wealth (or capital) to ensure political or systemic stability, while simultaneously facing a penalty for failing to meet a specified growth target. Because private sectors are typically more productive but also accumulate wealth faster, meeting high growth targets requires relaxing controls on private activity. However, this relaxation increases the private wealth share, eventually hitting the control constraint and forcing the planner to retrench. This cycle of expansion and retrenchment generates endogenous fluctuations in growth and policy stance, even in the absence of exogenous shocks.

We motivate this theoretical framework with stylized facts from China and the United States. First, we document the policy pendulum in China by analyzing the spread between private and state-owned enterprise (SOE) investment growth from 1981 to 2024. We identify distinct cycles corresponding to major policy shifts, such as the pro-market reforms of the 1980s and the post-2008 resurgence of state-led investment. Second, we examine the role of GDP growth targets. We show that the relationship between announced targets and realized growth has evolved into two regimes: a pre-2011 "high-growth" era where targets served as soft lower bounds, and a post-2011 "New Normal" where targets function as binding constraints, leading to managed deceleration. We highlight the risks of forcibly maintaining high growth targets through regional case studies of Liaoning and Inner Mongolia, where unsustainable growth pursuits led to sharp economic contractions. Third, we provide cross-country evidence from the U.S., showing that financial deregulation often coincides with credit expansions and asset price bubbles, while crises trigger swings back toward tighter control, illustrating the universal trade-off between efficiency and stability.

Our quantitative analysis demonstrates that the pendulum model can replicate the key features of China's growth experience. Calibrated to match historical data, the model generates endogenous

cycles of high growth accompanied by rising private wealth shares, followed by periods of slower growth and state consolidation. We further explore the implications of the model for policy design, showing that aggressive growth targets increase both mean growth and volatility, creating a mean-volatility trade-off. Importantly, we find that a "soft landing"—a smooth transition to lower growth—is achievable if the planner gradually lowers the growth mandate before the control constraint becomes binding, thereby avoiding the sharp corrections associated with defending unsustainable targets.

To establish the robustness of our mechanism, we extend the analysis to a standard neoclassical growth model. We show that the introduction of a growth mandate into a textbook framework with diminishing returns also generates investment cycles. In this setting, the state variable is aggregate capital rather than private wealth share, but the logic remains identical: defending a growth target requires higher investment when capital is abundant, which accumulates capital and makes future target defense increasingly costly, leading to cycles of over- and under-investment.

Related literature. This paper contributes to the literature on political economy and macroeconomic policy by providing a unified framework for understanding policy cycles in both authoritarian and democratic contexts. Existing studies often attribute policy volatility to exogenous shocks or political turnover (Nordhaus, 1975; Hibbs, 1977; Alesina, 1987; Rogoff, 1990; Persson and Tabellini, 1990). In contrast, we highlight the endogenous nature of these swings driven by the structural tension between growth and control. By linking theoretical mechanisms to empirical evidence from China and the U.S., we offer new insights into the dynamics of economic governance and the challenges of managing long-term growth transitions.

Our paper is also related to the literature on the theory of endogenous cycles. See, e.g., Benhabib and Nishimura (1979); Cass and Shell (1983); Woodford (1991); Benhabib and Farmer (1994); Boldrin and Rustichini (1994); Christiano and Harrison (1999); Beaudry et al. (2020). We differ from these studies by emphasizing that cycles are driven by endogenous swings in government policy, and the emergence of cycles in our environment is robust across almost the entire parameter space.

This paper is most closely related to studies on policy cycles in China (Lin et al., 1996; Naughton, 1996; Brandt and Zhu, 2000; Wang, 2021; Cheremukhin et al., 2024; Xu et al., 2025). Our work complements this literature by providing a unified framework that features perpetual cycles that do not converge to a steady state and quantifying its role in shaping China's economic fluctuations and growth. We also highlight a trade-off between volatility and growth rate faced by the policymaker.

The remainder of the paper is organized as follows. Section 2 presents the motivating evidence from China and the U.S. Section 3 develops the theoretical model. Section 4 provides quantitative explorations of the Chinese case. Section 5 extends the analysis to a neoclassical framework. Section 6 concludes.

2. Motivating Evidence

This section documents the stylized facts that motivate our theoretical framework next. We present empirical evidence regarding the cyclicality of economic policy—oscillating between regimes of state control and market liberalization—and its interaction with growth targets. While our primary focus is on the Chinese economy, we also provide evidence from the United States to illustrate that similar trade-offs between efficiency and stability manifest across different institutional settings.

2.1 The Policy Pendulum in China

We first document the presence of distinct policy cycles in China, characterized by alternating phases of "pro-private" and "pro-state" interventions. Figure 1 overlays three series for 1981–2024: the spread between the growth rate of private fixed asset investment and that of state-owned enterprises (SOEs) (shaded area), the aggregate GDP growth rate (gray bars), and a timeline of major policy changes and reforms (labeled boxes). The shaded regions indicate the direction of the spread: red areas denote periods where private investment outpaced SOE investment (a proxy for pro-private regimes), while blue areas denote the reverse. The timeline highlights major policy shifts: episodes marked in red correspond to pro-private reforms, while those in green represent pro-state interventions.

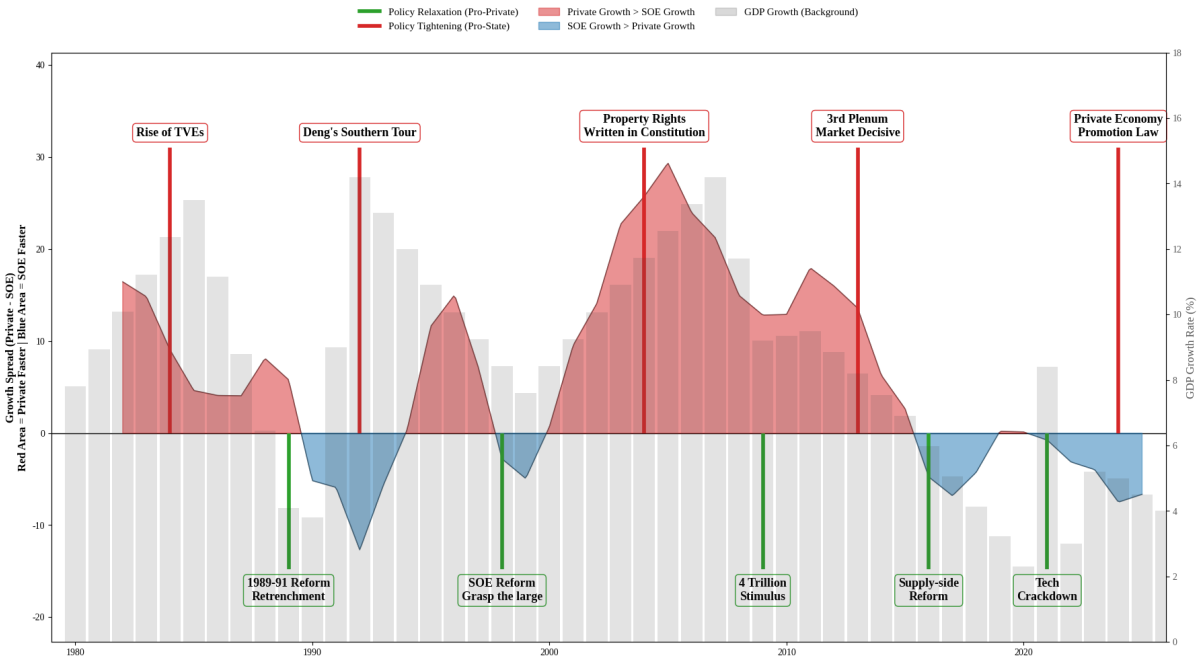
Notably, this figure reveals a "pendulum" pattern driven by major policy shifts:

- **Rise of TVEs (1980s)** The early reform era saw the explosion of Township and Village Enterprises (TVEs). Many private entities operated as "Red Hat" firms—nominally collective to obtain political protection but effectively private—leading to a surge in non-state activity.
- **Retrenchment (1989–1991)** Following the political turmoil of the late 1980s, the government implemented a "governance and rectification" program. Credit tightening and a shift in the ideological climate resulted in a sharp contraction of the private-SOE investment spread. (the first blue trough).
- **The Pro-Market Boom (1992–2000s)** Deng Xiaoping's "Southern Tour" in 1992 reinvigorated the reform agenda. This was followed by the landmark inclusion of "inviolability of private property" into the Constitution in 2004. These policies cemented confidence, leading to a decade-long dominance of private sector growth.[12]
- **The Reversal (Post-2008)** The pendulum swung back following the Global Financial Crisis. The "Four Trillion Yuan" stimulus was channeled primarily through SOEs. Despite the 2013 Third

[1]We classify "Grasp the Large" as state consolidation. While it privatized small firms, its strategic intent was to strengthen state control over key strategic sectors ("National Champions"), thereby enhancing the state's aggregate capacity over the economy's "commanding heights" despite reducing the total number of firms.

[2]The reversion to relatively higher SOE investment growth circa 1998 coincides with the Asian Financial Crisis. To maintain the growth mandate and defend the currency peg, the government launched a massive fiscal stimulus program centered on infrastructure investment primarily executed by the state sector. The data capture a counter-cyclical expansion of SOEs aimed at macroeconomic stabilization rather than ideological retrenchment.

Figure 1: Policy Cycles in China's Pro-SOE and Pro-Private Policies



Notes: The figure displays the spread between the growth rate of private fixed asset investment and that of state-owned enterprises (SOEs) (left axis, shaded areas), alongside the real GDP growth rate (background bars). Data are from the National Bureau of Statistics (NBS). Due to statistical standard changes, the series is constructed in two phases: (1) 1981–2003 uses "Non-State" vs. "State-owned Units" investment; (2) 2004–2024 uses officially published "Private" vs. "State-owned and State-holding" investment, with a calculated transition for the 2004 bridge year. The spread series is smoothed using a three-year moving average. Red (blue) shaded areas indicate periods where private (state) investment growth exceeds state (private) investment growth. Text boxes denote major policy events. See footnote 1 and 2 for the classification of "Grasp the Large" and the blue trough circa 1998.

Plenum's declaration to let the "market play a decisive role," subsequent policies—such as the "Supply-Side Structural Reform" (which reduced capacity largely in private SMEs) and the recent crackdown on platform economy and technology sectors (e.g., Alibaba, Tencent) to prevent the "disorderly expansion of capital"—have reasserted state control.

- **Recent Stabilization (2025-)** Most recently, in response to downward pressure on growth, policy signals have shifted again toward supporting the private sector, exemplified by the drafting of the Private Economy Promotion Law.

This historical trajectory highlights a recurring oscillation: periods of deregulation and private sector expansion are often followed by periods of re-regulation and state consolidation.

2.2 Growth Mandates and the "Soft Landing"

A second key feature of the Chinese macroeconomy is the role of explicit GDP growth targets. Figure 2 compares the government's announced growth target against the realized GDP growth rate from 1985 to 2025.

Figure 2: China's GDP Growth Targets and the Realized Growth Rates



Notes: The solid line represents the official real GDP growth target announced by the central government in the annual Government Work Report. The bars represent the realized annual real GDP growth rate reported by the NBS. The figure illustrates the evolution from a "high-growth" regime (pre-2011), where targets acted as soft lower bounds, to a "New Normal" regime (post-2011), where targets function as binding constraints characterized by managed deceleration and high compliance.

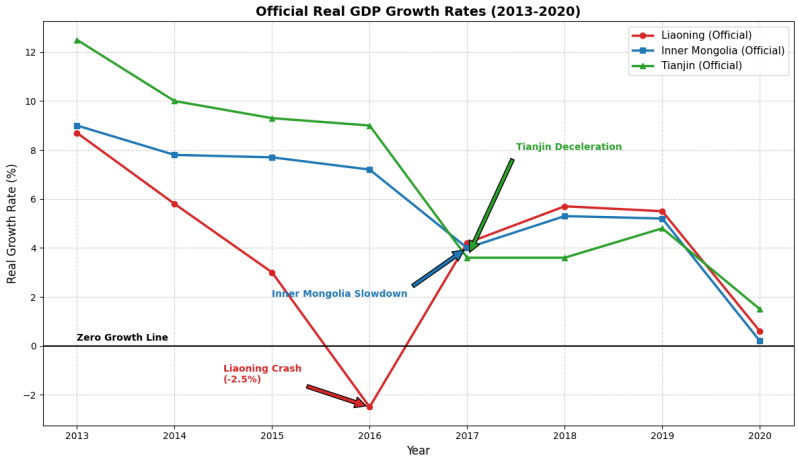
The relationship between the target and realized growth has evolved into two distinct regimes. In the high-growth era (prior to 2011), the growth target functioned as a soft lower bound, with realized growth frequently exceeding the target by substantial margins (e.g., during the "Keep Eight" era). However, as the economy entered a phase of slower potential growth (the "New Normal"), the target became a binding constraint. Realized growth in this latter period has tracked the target with high precision.

Furthermore, the data suggests a preference for managed deceleration. As shown in Figure 2, the official target has been gradually lowered from 8% to around 5%. Figure 3 provides regional evidence illustrating the potential risks associated with forcibly pursuing high growth rates. Provinces such as Liaoning and Inner Mongolia, which tried to maintain high growth rates despite deteriorating fundamentals, eventually experienced sharp economic contractions and/or statistical revisions. Such discontinuities do not appear at the national level, suggesting that policymakers adjust targets to smooth the transition.

2.3 Cross-Country Evidence: U.S. Financial Cycles

Finally, we provide suggestive evidence that the tension between control and efficiency is not unique to China. Figure 4 provides an analogous time-series depiction for the U.S. since 1980, plotting a measure of credit expansion/tightening (shaded area), a detrended S&P 500 index (dashed line), and major policy changes in financial regulation and deregulation (labeled), with the financial crisis marked

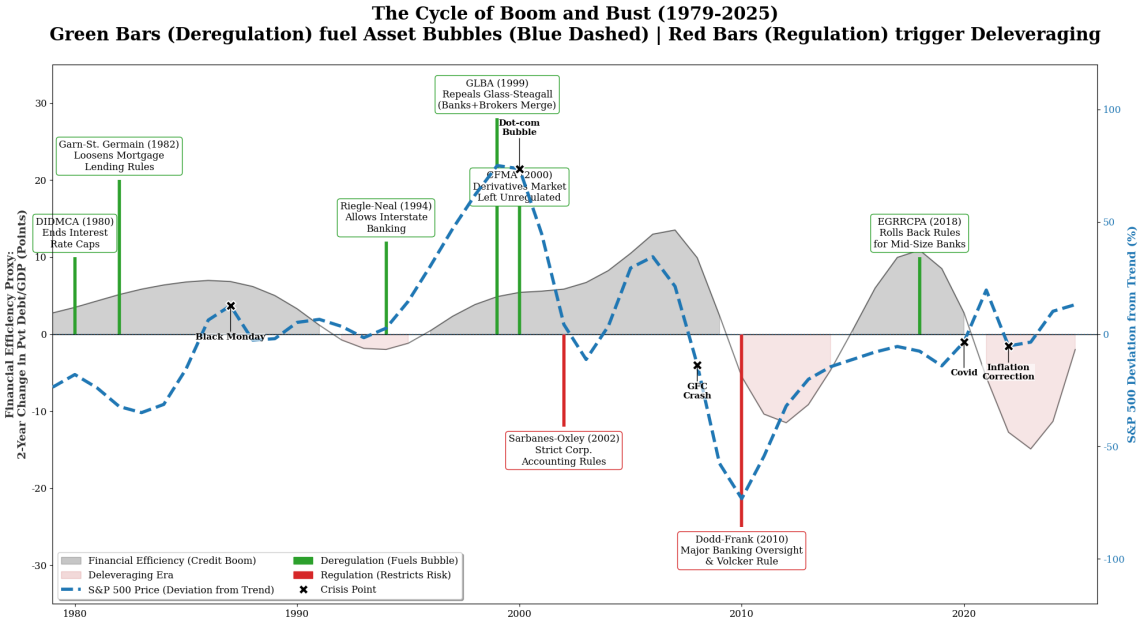
Figure 3: Regional Sudden Stops



Notes: The figure plots the nominal GDP growth rates for Liaoning and Inner Mongolia provinces, using data from Provincial Statistical Yearbooks. The sharp declines observed (statistical "sudden stops") illustrate the consequences of maintaining high growth that are inconsistent with deteriorating local economic fundamentals, eventually necessitating drastic corrections.

explicitly.

Figure 4: Policy Cycle in the US Financial Regulations and Deregulations



Notes: This figure illustrates the co-movement of financial cycles and regulatory regimes using data from the FRED and S&P Global. The blue dashed line plots the cyclical component of the S&P 500 index, calculated as the percentage deviation of real prices from a long-run log-linear trend. Background shaded regions represent the credit impulse, defined as the two-year change in the U.S. private non-financial sector debt-to-GDP ratio, where grey areas denote periods of credit expansion and red areas indicate phases of deleveraging. Vertical bars mark structural policy shifts, with green bars corresponding to deregulatory measures (e.g., the 1999 Gramm-Leach-Bliley Act) and red bars corresponding to regulatory tightening (e.g., the 2010 Dodd-Frank Act). Markers highlight major financial crises to demonstrate the synchronization between market turning points and the regulatory cycle.

The U.S. financial system exhibits a similar cycle. Periods of deregulation, such as the repeal of the Glass-Steagall Act (1999), often coincide with credit and asset price expansions. Although this may enhance economic efficiency, it also introduces potential bubbles in asset prices. Conversely, major crises typically trigger a swing back toward tighter control, exemplified by the Sarbanes-Oxley Act (2002) and the Dodd-Frank Act (2010). This suggests that the trade-off between efficiency and stability helps explain regulatory cycles across different institutional settings.

3. Theory

3.1 Environment and the Planner's Problem

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. The economy is populated by two representative households, a private household p and a state household s. The private household owns the wealth of private firms, and the state household owns the wealth of state firms. Let $a_{p,t}$ and $a_{s,t}$ denote their beginning-of-period wealth.

Households. The economy is populated by two representative households, a private household p and a state household s. Each household is forward looking and takes the return on its wealth as given. Household $j \in \{p, s\}$ accumulates wealth by solving

$$\max_{\{c_{j,t}, a_{j,t+1}\}_{t\geq 0}} \sum_{t=0}^{\infty} \beta_j^t \log c_{j,t}$$
$$s.t. \quad c_{j,t} + a_{j,t+1} = R_{j,t} a_{j,t}$$

where $a_{j,t}$ is the beginning-of-period wealth, β_j is the discount factor, and the return $R_{j,t}$ is taken as given by the household.

Production. Household $j \in \{p, s\}$ has access to a linear production technology through a representative firm also labeled as j. In particular, firm j produces the final good Y using capital according to

$$Y_{j,t} = z_j K_{j,t}$$

where z_j is the productivity of firm j. Household j does not have access to the other production technology. We assume that the private firm is more productive, i.e., $z_p > z_s$.

Policy instrument. The planner controls the amount of capital that the private household can invest in the private firm. In particular, the planner sets λ_t such that

$$K_{p,t} \leq \lambda_t a_{p,t}$$

In the baseline model, we assume $\lambda_t \in [0, 1]$. Therefore, the private household can only use part of its wealth to produce in the private firm. It lends the rest of its wealth to the state household through a competitive lending market, and the state household can use that to produce in the state firm. A higher λ_t means the private household can direct more of its wealth to the private firm.

The planner's problem. The planner has two targets in the economy. Denote by μ_t the private wealth share, i.e.,

$$\mu_t = \frac{a_{p,t}}{a_{p,t} + a_{s,t}}$$

and denote by g_t the output growth rate, i.e.,

$$g_t = \frac{Y_{p,t} + Y_{s,t}}{Y_{p,t-1} + Y_{s,t-1}}$$

The planner faces two objectives: containing the accumulation of private wealth and maintaining a sufficiently high growth rate. In addition to the general preference for higher growth, the planner also faces an explicit growth mandate. In particular, we assume the planner's objective function is

$$\max_{\{\lambda_t\}_{t\geq 0}} \sum_{t=0}^{\infty} \beta_g^t \left[\alpha \log(g_t) + (1-\alpha)\log(\bar{\mu} - \mu_{t+1}) - \chi \mathbb{I}(g_t < \bar{g})\right]$$

subject to the initial private wealth share μ_0, and the path of $\{\mu_{t+1}, g_t\}_{t\geq 0}$ induced by its policy path $\{\lambda_t\}_{t\geq 0}$ in the decentralized economy. In this objective function, $\alpha \in (0, 1)$ captures the relative weight the planner puts on growth and on the private wealth share, $\bar{\mu} \in (0, 1)$ is an upper bound for the private wealth share and captures the state control over private wealth, $\bar{g}$ is the growth mandate that the planner faces, and χ is the per-period punishment if the planner fails to meet the growth mandate, indicated by the indicator function. However, unlike the growth mandate, the cap on the private wealth share is a hard constraint that the planner can never violate, which is a consequence of the log formulation.

3.2 The Decentralized Economy for a Given Policy Path

We first characterize the decentralized economy for any given path of the policy instrument. In period t, given the capital cap λ_t and the beginning-of-period wealth $a_{p,t}, a_{s,t}$ of the two households, the private firm uses $K_{p,t} = \lambda_t a_{p,t}$ in production, while the state firm uses $K_{s,t} = a_{s,t} + (1 - \lambda_t)a_{p,t}$ in production. Given this capital allocation, total output is

$$Y_t = z_p K_{p,t} + z_s K_{s,t} = z_s K_t + \lambda_t (z_p - z_s) a_{p,t}$$

where $K_t = K_{p,t} + K_{s,t}$ is the total capital stock in period t.

Since household s puts all of its capital in the state firm, its capital return is

$$R_s = z_s$$

For household p, a fraction λ_t of its wealth earns return z_p, and the rest earns return z_s. So, its effective wealth return is

$$R_p(\lambda_t) = \lambda_t z_p + (1 - \lambda_t) z_s$$

Log utility yields a simple consumption-saving rule for household $j \in \{p, s\}$: it saves a fraction β_j of its end-of-period wealth, as stated in the next lemma.

Lemma 3.1. *Given the capital cap λ_t and the beginning-of-period wealth $a_{p,t}, a_{s,t}$, the state and private wealth evolve as*

$$a_{p,t+1} = \beta_p R_p(\lambda_t) a_{p,t}, \quad a_{s,t+1} = \beta_s R_s a_{s,t}$$

With Lemma 3.1, we can characterize the planner's two targets, private wealth share in the next period μ_{t+1}, and the output growth rate g_t. In particular, we can write them both as functions of the current private wealth share μ_t, and the policy λ_t.

Lemma 3.2. *Given the current private wealth share μ_t and policy λ_t, the private wealth share next period and the output growth rate are*

$$\mu_{t+1} = \mathcal{M}(\mu_t, \lambda_t) := \frac{\mu_t \beta_p R_p(\lambda_t)}{(1 - \mu_t)\beta_s R_s + \mu_t \beta_p R_p(\lambda_t)}$$
$$g_t = \mathcal{G}(\mu_t, \lambda_t) := \frac{z_s + (z_p - z_s)\lambda_t \mu_t}{\frac{1}{\beta_s} + \frac{\beta_s - \beta_p}{\beta_s \beta_p} \mu_t}$$

Moreover, μ_{t+1} increases with both μ_t and λ_t, and g_t increases with λ_t.

Lemma 3.2 already hints some policy tradeoff the planner faces. On the one hand, a looser private capital restriction (λ_t higher) raises the private wealth share next period, as the private household gets a higher capital return. This is not desirable to the planner. On the other hand, a higher λ_t leads to a higher output growth rate, as more capital is directed to the more efficient private firm. The planner likes it, and it helps to meet the growth mandate.

3.3 The Planner's Problem: The Myopic Case

Lemma 3.2 also implies that the current private wealth share μ_t is the only state variable that the planner needs to keep track of. In other words, the planner's problem has the following recursive formulation.

Lemma 3.3. *The planner's problem has the following recursive formulation:*

$$V(\mu_t) = \max_{\lambda_t \in [0,1]} \alpha \log(g_t) + (1 - \alpha) \log(\bar{\mu} - \mu_{t+1}) - \chi \mathbb{I}(g_t < \bar{g}) + \beta_g V(\mu_{t+1})$$

$$s.t.\ \ \mu_{t+1} = \mathcal{M}(\mu_t, \lambda_t)$$
$$g_t = \mathcal{G}(\mu_t, \lambda_t)$$

In general, the continuation value $\beta_g V(\mu_{t+1})$ complicates the problem. To study the planner's problem in a more transparent way, in this subsection, we assume the planner is myopic, i.e., $\beta_g = 0$. This effectively shuts down the dynamic consideration and isolates the tradeoff between the output growth rate and the state control over the private wealth share, which is the planner's primary concern. As will become clear shortly, this special case can already yield interesting dynamics. We solve for the full dynamic model numerically in the quantitative part.

In the myopic case, the time t planner chooses λ_t only to maximize the current-period payoff, taking the current private wealth share μ_t as given. To simplify notation, in this subsection we drop the time subscript and write

$$U(\mu, \lambda) = \alpha \log \mathcal{G}(\mu, \lambda) + (1 - \alpha) \log (\bar{\mu} - \mathcal{M}(\mu, \lambda)) - \chi \mathbb{I}(\mathcal{G}(\mu, \lambda) < \bar{g})$$

Denote by $\Lambda(\mu)$ the feasible set when current private wealth share is μ, which is

$$\Lambda(\mu) = \{\lambda \in [0, 1] : \mathcal{M}(\mu, \lambda) < \bar{\mu}\}$$

Denote by $\lambda^m(\mu)$ the optimal myopic policy rule, and $\Phi(\mu)$ the resulting private wealth share next period following that policy,[3] i.e.,

$$\lambda^m(\mu) = \arg \max_{\lambda \in \Lambda(\mu)} U(\mu, \lambda), \qquad \Phi(\mu) = \mathcal{M}(\mu, \lambda^m(\mu)).$$

The function Φ characterizes the law of motion of the private wealth share induced by the myopic planner. The rest of this subsection builds the mechanism in three steps. First, we identify the policy that keeps the private wealth share constant. Second, we show that a nontrivial growth mandate forces the planner to choose above that policy whenever it wants to satisfy the mandate. Third, we show why this growth mandate is not sustainable all the way to the control cap on the private wealth share. The policy pendulum is the result of putting these two forces together.

We impose two restrictions that make the tradeoff meaningful. The first one says that the private household does not automatically dominate the economy unless policy gives it enough access to the private technology. The second one says that the growth target is high enough to bind, but not so high that it is mechanically impossible.

Assumption 1. *Private firms are more productive, but private households are more impatient:*

$$\beta_p z_s < \beta_s z_s < \beta_p z_p.$$

[3] We assume that whenever the planner is indifferent between meeting the mandate and missing it, the planner chooses to meet the mandate. Under this tie-breaking rule, one can show that the optimal policy is unique for all $\mu \in [0, \bar{\mu})$.

This assumption makes the policy instrument nontrivial. If private wealth is operated by the state firm, private wealth grows more slowly than state wealth. If private wealth is operated by the private firm, private wealth grows faster than state wealth. Therefore, there is an interior policy that keeps the private wealth share constant.

Assumption 2. *The growth mandate is nontrivial:*

$$\beta_s z_s < \bar{g} < \mathcal{G}(\bar{\mu}, 1).$$

The lower bound says that the planner cannot satisfy the mandate while keeping the private wealth share constant. The upper bound rules out the case in which the mandate is never attainable even under the loosest private capital restriction. Thus the mandate is feasible in principle, but incompatible with a stationary private wealth share.

Thresholds. The useful way to read the model is not to start from the planner's objective, but from three policy cutoffs. Each cutoff answers a different question. How much private access keeps the private wealth share unchanged? How much private access is needed to hit the growth mandate? How much private access would push the economy exactly to the control cap? The relative location of these cutoffs tells us whether the planner is pushed toward expansion, i.e., meeting the growth mandate, or retrenchment, i.e., suppressing the private wealth share.

The first cutoff is the stationary-share policy λ^{ss}. It solves

$$\beta_p R_p(\lambda^{ss}) = \beta_s R_s$$

Since $R_p(\lambda) = z_s + \lambda(z_p - z_s)$ and $R_s = z_s$, Lemma 3.1 implies that policies above λ^{ss} raise the private wealth share, and policies below λ^{ss} reduce it.

The second cutoff is the growth-target policy $\lambda^g(\mu)$. It is the lowest policy that reaches the growth mandate:

$$\mathcal{G}(\mu, \lambda^g(\mu)) = \bar{g}$$

When $\lambda^g(\mu) \in [0, 1]$, the mandate can be met at state μ. When $\lambda^g(\mu) > 1$, the current private sector is too small for the mandate to be attainable under the baseline cap.

The third cutoff is the control-cap policy $\lambda^c(\mu)$. It solves

$$\mathcal{M}(\mu, \lambda^c(\mu)) = \bar{\mu}$$

For $\mu \in (0, \bar{\mu})$, the feasible set can be written as $\Lambda(\mu) = [0, 1] \cap [0, \lambda^c(\mu))$ whenever the right-hand side is nonempty.

The no-mandate benchmark. Before using the mandate, it is useful to look at the policy the planner would choose if missing $\bar{g}$ carried no fixed punishment. This benchmark is not meant to be the final

policy rule. It gives us the smooth tradeoff that the mandate will later distort. Define

$$u^0(\mu, \lambda) = \alpha \log \mathcal{G}(\mu, \lambda) + (1 - \alpha) \log (\bar{\mu} - \mathcal{M}(\mu, \lambda))$$

and

$$\lambda^0(\mu) = \arg \max_{\lambda \in \Lambda(\mu)} u^0(\mu, \lambda)$$

The smooth benchmark need not generate cycles. The policy rule changes continuously with the state, and the economy may converge to a steady state where the marginal growth gain from private access equals the marginal control cost.The mandate changes the problem qualitatively because it gives the planner a discrete reward for crossing the growth threshold.

The next lemma links the cutoffs to the law of motion of the private wealth share. At λ^{ss}, the planner exactly holds the private wealth share fixed. This makes λ^{ss} the natural dividing line: policies above it move the economy toward private expansion, while policies below it move the economy back toward state control.

Lemma 3.4. *For any $\mu \in (0, 1)$,*

$$\mathcal{M}(\mu, \lambda^{ss}) = \mu, \qquad \mathcal{G}(\mu, \lambda^{ss}) = \beta_s z_s.$$

Proof. At λ^{ss}, private and state wealth grow at the same rate: $\beta_p R_p(\lambda^{ss}) = \beta_s R_s = \beta_s z_s$. Substituting this equality into the law of motion for the private wealth share gives $\mathcal{M}(\mu, \lambda^{ss}) = \mu$. The same equality implies that aggregate capital, and therefore output, grows at rate $\beta_s z_s$. □

We can now state the central incompatibility. If the mandate were no higher than the growth rate generated at λ^{ss}, the planner could satisfy growth and keep the private wealth share constant at the same time. Assumption 2 rules this out. Meeting the mandate necessarily means choosing a policy that raises the private wealth share next period.

Proposition 1. *Under Assumptions 1 and 2, any feasible policy that meets the growth mandate must lie above the stationary-share policy. In particular, whenever $\lambda^g(\mu) \in \Lambda(\mu)$,*

$$\lambda^g(\mu) > \lambda^{ss}, \qquad \mathcal{M}(\mu, \lambda^g(\mu)) > \mu.$$

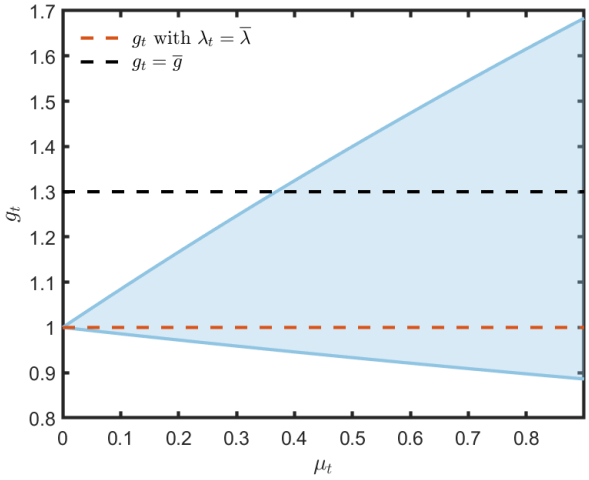
Proof. By Lemma 3.4, the stationary-share policy delivers growth $\beta_s z_s$. Assumption 2 implies $\bar{g} > \beta_s z_s$. Since $\mathcal{G}(\mu, \lambda)$ is strictly increasing in λ for $\mu > 0$, meeting the mandate requires $\lambda^g(\mu) > \lambda^{ss}$. Since $\mathcal{M}(\mu, \lambda)$ is also strictly increasing in λ, the target-meeting policy raises the next private wealth share above μ. □

This is the core tradeoff. The planner has only one instrument. The same relaxation of the private capital restriction raises current growth but also raises the future private wealth share. Figure 5 illustrates Proposition 1 in a more transparent way. The shaded area corresponds to the possible growth

rates for each of private wealth share μ_t. The upper bound is obtained by setting $\lambda_t = 1$ and the lower bound is obtained by setting $\lambda_t = 0$. With a higher share of private wealth, a higher growth rate level can be potentially achieved. When setting the financial constraint $\lambda_t = \overline{\lambda}$, the level that ensures a constant share of private wealth, the growth rate is equal to $\beta_s z_s$. This level corresponds to the red dashed line in Figure 5, which is generally lower than the economic growth potential.

Consider the case in which the planner wants to achieve a growth rate $\overline{g} > \beta_s z_s$. This is possible only if μ_t is high enough and the required λ_t is above $\overline{\lambda}$. However, setting $\lambda_t > \overline{\lambda}$, it will eventually lead to the scenario in which μ_t approaches 1. This is the challenge that the planner faces: It is impossible to maintain a relatively high growth rate while admitting a relatively high share of state wealth.

Figure 5: Growth Rates



The incompatibility result says what must happen if the planner meets the mandate. It does not yet say when the planner is willing to do so. The answer depends on the size of the mandate punishment. When χ is small, the planner may prefer to miss the target and stay close to the smooth no-mandate policy. When χ is large, avoiding the fixed loss can justify a discrete upward move in λ. The next result formalizes this local expansion force.

Proposition 2. *Suppose $\mathcal{G}(\mu, \lambda^0(\mu)) < \bar{g}$ and $\lambda^g(\mu) \in \Lambda(\mu)$. Then there exists a finite threshold $\chi^*(\mu)$ such that the myopic planner chooses a target-meeting policy whenever $\chi \geq \chi^*(\mu)$.*

Proof. The best payoff from missing the mandate is $u^0(\mu, \lambda^0(\mu)) - \chi$. The best payoff from satisfying the mandate is

$$W^g(\mu) = \sup_{\lambda \in \Lambda(\mu),\ \mathcal{G}(\mu,\lambda) \geq \bar{g}} u^0(\mu, \lambda).$$

Thus the planner satisfies the mandate whenever

$$\chi \geq u^0(\mu, \lambda^0(\mu)) - W^g(\mu) \equiv \chi^*(\mu).$$

□

Combining Propositions 1 and 2, the mandate creates an expansionary region in which $\lambda^m(\mu) > \lambda^{ss}$ and therefore $\Phi(\mu) > \mu$. In a phase diagram, the law of motion lies above the 45-degree line in this region.

This expansion matters for the dynamics because it can remove a steady state. In the smooth benchmark, a fixed point is simply a place where the policy rule crosses the 45-degree line. The mandate can lift the policy rule above that line around the same state. When this happens, the economy no longer has a local resting point there. Suppose the no-mandate law of motion has an interior fixed point μ^0 with $\mathcal{M}(\mu^0, \lambda^0(\mu^0)) = \mu^0$ and $\mathcal{G}(\mu^0, \lambda^0(\mu^0)) < \bar{g}$. If χ is large enough that the planner jumps to a mandate-satisfying policy around μ^0, then $\Phi(\mu) > \mu$ in that neighborhood. The crossing with the 45-degree line that existed in the smooth benchmark disappears. For some parameter values, a high enough χ expands this upward-drift region so much that Φ has no fixed point on the relevant feasible interval.

However, the expansionary force has a natural limit. The mandate penalty is a fixed number, while the control cost becomes arbitrarily large as the next private wealth share approaches $\bar{\mu}$. Thus even a planner facing a large mandate penalty eventually becomes unwilling to buy growth by moving too close to the cap.

Proposition 3. *For any finite χ, there exists $\varepsilon > 0$ such that*

$$\lambda^m(\mu) < \lambda^{ss} \quad \textit{and} \quad \Phi(\mu) < \mu \quad \textit{for all } \mu \in (\bar{\mu} - \varepsilon, \bar{\mu}).$$

Proof. Under Assumption 1, the policy $\lambda = 0$ keeps a positive distance from the cap even as μ is arbitrarily close to $\bar{\mu}$, because $\mathcal{M}(\bar{\mu}, 0) < \bar{\mu}$. Hence the payoff from $\lambda = 0$ remains bounded from below near the cap. By contrast, for any $\lambda \geq \lambda^{ss}$, Lemma 3.4 and the monotonicity of $\mathcal{M}$ imply $\mathcal{M}(\mu, \lambda) \geq \mu$. Therefore $\bar{\mu} - \mathcal{M}(\mu, \lambda) \leq \bar{\mu} - \mu$, and the control term goes to $-\infty$ as $\mu \uparrow \bar{\mu}$. Since the growth benefit and the mandate penalty are bounded, all policies with $\lambda \geq \lambda^{ss}$ are dominated by $\lambda = 0$ sufficiently close to the cap. Thus $\lambda^m(\mu) < \lambda^{ss}$ and $\Phi(\mu) < \mu$ in a neighborhood of $\bar{\mu}$. □

The expansionary force generated by the mandate is therefore state dependent. It is strong when the private wealth share is low or moderate, because the planner can raise growth without immediately endangering control. Near the cap, control dominates growth. The planner retrenches by lowering λ and accepting weaker current growth in order to restore room below $\bar{\mu}$. Notice Proposition 3 applies to the no-mandate benchmark as well ($\chi = 0$). Namely, there exists some ε^0 such that $\lambda^0(\mu) < \lambda^{ss}$ for all $\mu \in (\bar{\mu} - \varepsilon^0, \bar{\mu})$. This property will be useful next.

We now have the two pieces needed for a pendulum. For a sufficiently large mandate penalty, there is an interior region in which the planner expands private access and pushes μ upward. For any finite mandate penalty, there is a region near the control cap in which the planner retrenches and pushes μ downward. The final result puts these two statements together.

Proposition 4. *Suppose there exists* $\mu^e \in (\bar{\mu} - \varepsilon^0, \bar{\mu})$ *and* $\lambda^e \in \Lambda(\mu^e)$ *such that*

$$\mathcal{G}(\mu^e, \lambda^e) > \bar{g}, \qquad \mathcal{G}(\mu^e, \lambda^0(\mu^e)) < \bar{g}$$

Then, there exists some $\bar{\chi}$ *such that, for every finite* $\chi \geq \bar{\chi}$*, there exist* $\mu_L < \mu_M < \mu_H(\chi) < \bar{\mu}$ *such that*

$$\lambda^m(\mu) > \lambda^{ss}, \Phi(\mu) > \mu, \quad \forall \mu \in (\mu_L, \mu_M),$$

$$\lambda^m(\mu) \neq \lambda^{ss}, \Phi(\mu) \neq \mu, \quad \forall \mu \in (\mu_M, \mu_H(\chi)),$$

and

$$\lambda^m(\mu) < \lambda^{ss}, \Phi(\mu) < \mu, \quad \forall \mu \in (\mu_H(\chi), \bar{\mu}).$$

The condition says that there is at least one interior state where the growth mandate is genuinely relevant and causes a conflict between growth and control. It is attainable, because some feasible policy λ^e reaches the target with slack and without violating the control cap. It is also not automatically satisfied, because the smooth no-mandate policy $\lambda^0(\mu^e)$ would miss the target. By continuity, the same economic situation holds in a small interval around μ^e. On this interval, a large enough mandate penalty makes the planner cross the growth threshold. In addition, absent the growth mandate, this point lies in a region where the planner wants to reduce the private wealth share.

Proof. By continuity, there exists $\eta > 0$ such that $I = [\mu^e - \eta, \mu^e + \eta] \subset (\bar{\mu} - \varepsilon^0, \bar{\mu})$ and, for every $\mu \in I$, the mandate is feasible with slack and $\mathcal{G}(\mu, \lambda^0(\mu)) < \bar{g}$. For each $\mu \in I$, Proposition 2 gives a finite threshold $\chi^*(\mu)$ above which the planner chooses a target-meeting policy. Since I is compact, the payoff objects are continuous away from the cap, and the mandate is feasible with slack throughout I, $\bar{\chi} = \sup_{\mu \in I} \chi^*(\mu)$ is finite. For every $\chi \geq \bar{\chi}$, the planner meets the mandate throughout I. Proposition 1 then implies $\lambda^m(\mu) \geq \lambda^g(\mu) > \lambda^{ss}$, and Lemma 3.4 implies $\Phi(\mu) > \mu$, for all $\mu \in I$. Set $\mu_L = \mu^e - \eta$ and $\mu_M = \mu^e + \eta$.

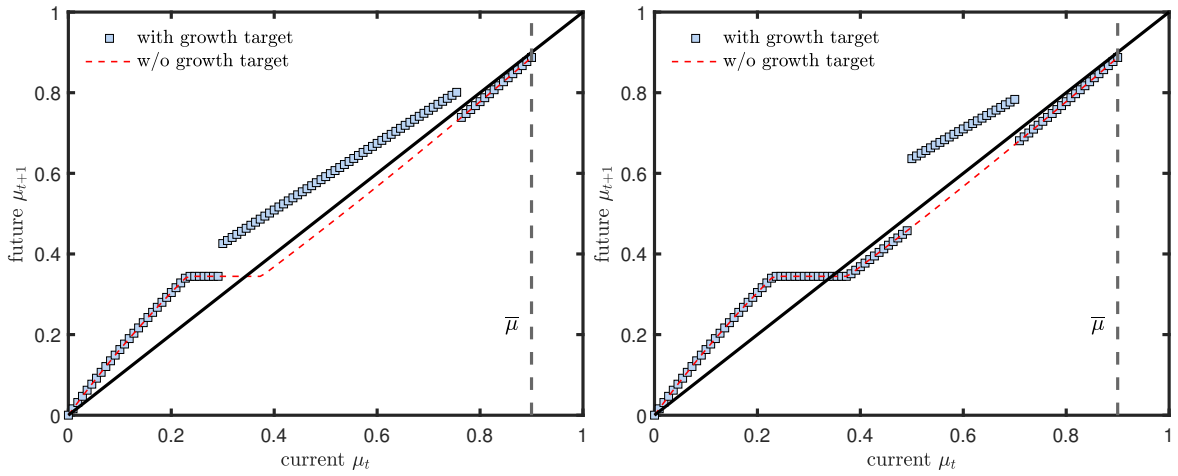
Then, for all $\mu \in (\mu_M, \mu_H(\chi))$, we still have $\mu \in (\bar{\mu} - \varepsilon^0, \bar{\mu})$, so the no-mandate policy satisfies $\lambda^0(\mu) < \lambda^{ss}$. If the planner chooses to miss the growth mandate, the penalty term is fixed and the best missing policy is the no-mandate maximizer $\lambda^0(\mu)$. This policy indeed misses the mandate, because $\lambda^0(\mu) < \lambda^{ss}$, $\mathcal{G}$ is increasing in λ, and $\mathcal{G}(\mu, \lambda^{ss}) = \beta_s z_s < \bar{g}$. Hence giving up the mandate implies $\lambda^m(\mu) = \lambda^0(\mu) < \lambda^{ss}$. If instead the planner chooses to meet the target, Proposition 1 implies $\lambda^m(\mu) > \lambda^{ss}$. In neither case does the planner set $\lambda^m(\mu)$ to λ^{ss}.

It remains to show that retrenchment must occur near the cap. For any finite $\chi \geq \bar{\chi}$, this follows directly from Proposition 3. It gives an $\varepsilon(\chi) > 0$ such that retrenchment holds for all $\mu \in (\bar{\mu} - \varepsilon(\chi), \bar{\mu})$. Choose $\mu_H(\chi) \in (\max\{\mu_M, \bar{\mu} - \varepsilon(\chi)\}, \bar{\mu})$. Then $\lambda^m(\mu) < \lambda^{ss}$ and $\Phi(\mu) < \mu$ for all $\mu \in (\mu_H(\chi), \bar{\mu})$. □

This is the policy pendulum. The mandate creates an expansionary interval in which $\Phi(\mu) > \mu$, while the control motive creates a retrenchment interval near the cap in which $\Phi(\mu) < \mu$. Together,

these statements rule out an interior resting point where $\lambda^m(\mu) = \lambda^{ss}$ for $\mu \in (\mu_L, \bar{\mu})$, and the two boundary points are not absorbing states either. By the feasibility constraint, it must be that $\Phi(\mu) < \bar{\mu}$. Otherwise, the planner will receive infinite punishment because of the log functional form. Therefore, if $\Phi(\mu)$ remains in $(\mu_L, \bar{\mu})$ for all $\mu \in (\mu_M, \bar{\mu})$, the economy must oscillates in this interval. A sufficient condition for this to be true is $\mathcal{M}(\mu_M, 0) > \mu_L$. [4] A characterization based on primitives is not easy to find, but by following the idea of this proposition to choose an appropriate χ, we can almost always generate oscillation in this economy numerically.

Figure 6: Evolution of Wealth Share with Growth Rate Mandate



Note: The left panel displays the law of motion with a relative low level of $\overline{g}$ and there does not exist a steady state. The right panel displays the law of motion with a relative high level of $\overline{g}$ and there exists a steady state.

Economically, a pro-private policy relaxes the capital restriction, raises current growth, and helps the planner meet the mandate. But the same policy raises the future private wealth share and gradually tightens the control constraint. Once the economy gets close enough to the control cap, the planner reverses course and restricts private access. Retrenchment restores control, but it also lowers growth and eventually recreates the pressure to expand.

Figure 6 displays the possible dynamics. The red dashed line in both panels is the law of motion of the private wealth share in the no-mandate benchmark. Under Assumption 1, the no-mandate benchmark has a unique interior steady state where $\lambda = \lambda^{ss}$. With a growth mandate, in the left panel, there is no interior steady state. When starting from a relatively low level of μ_t, the economy first experiences a monotonic increase of μ_t, following an endogenous cycle.

In the right panel, though there is an interior steady state, there are two possible regimes. If the economy starts with a relatively low level of μ_t, it will converge to μ^* monotonically. However, if the economy starts from a relatively high level of μ_t, an endogenous cycle will emerge and the economy never converges.

[4] Because on the one hand, any feasible choice makes sure $\mathcal{M}(\mu, \lambda) < \bar{\mu}$; on the other hand, $\mathcal{M}(\mu, \lambda)$ increases with both μ and λ , and therefore, $\mathcal{M}(\mu, \lambda) \geq \mathcal{M}(\mu_M, 0) > \mu_L$ for all $\mu \in (\mu_M, \bar{\mu}), \lambda \geq 0$.

To connect with China's economic history, one may interpret that before China's economic reform, the economy was stuck at the low steady state μ^*. After the economic reform, the economy started to experience an expansion of private firms' wealth share. However, once this share is high enough, the economy starts to alternate between favoring the state firms and private firms.

4. Quantitative Exploration

In the last section, we theoretically analyze the policy tradeoff of a myopic planner: a high growth mandate pushes the planner toward pro-private policies, while the control motive eventually forces retrenchment once the private wealth share becomes too high. This generates the growth pendulum: periods of private-sector expansion are followed by episodes of renewed state control. This section goes back to the forward-looking planner and asks whether a disciplined parameterization can generate the observed pattern of the Chinese economy, i.e., persistent growth management and recurrent swings between private-sector expansion and retrenchment.

Table 1: Parameterization

Param.	Value	Description
Fundamental Parameters		
β_p	0.8	Private sector discount factor
β_s	0.95	State sector discount factor
β	0.95	Planner discount factor
Z_s	$1.045/\beta_s$	Public sector TFP
Z_p	$1.25Z_s$	Private sector TFP
Policy Parameters		
a_0	0.8	Private sector capital share cap
α	0.60	Weights on growth rate in planner objective
χ	2.00	Punishment for not meeting growth mandate
$g_{mandate}$	1.06	Growth mandate

Notes: This table summarizes the model parameterization based on the fundamental and policy parameters used in the baseline calibration.

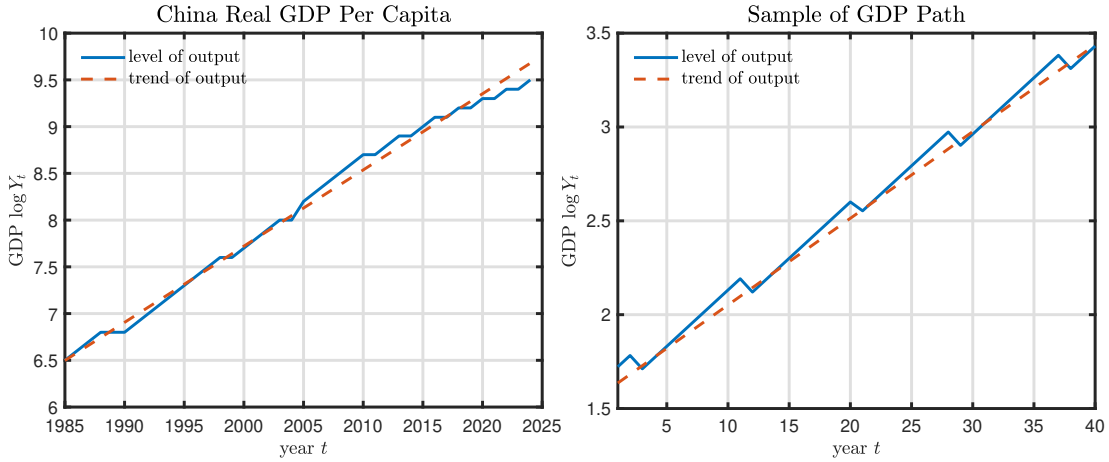
The baseline calibration is summarized in Table 1. We set the private-sector discount factor to 0.8, the public-sector discount factor and the planner's discount factor to 0.95. Public-sector productivity is normalized so that the public technology delivers gross growth 1.045 if public sector owns all the wealth, and private-sector productivity is 25% higher than public-sector productivity. The policy parameters are a private wealth-share cap of 0.8, a growth weight of 0.60, a mandate penalty of 2, and a baseline gross growth mandate of 1.06. In the simulations below, the planner is allowed a small amount of public lending to the private technology, with the policy instrument bounded above by

1.05 (i.e., $\lambda < 1.05$). This extension gives the planner somewhat more room to support growth, but it does not remove the central tension between growth and control.

4.1 Endogenous Growth Cycles

We first compare the model-implied output path with the Chinese GDP experience documented in the motivating evidence. The model is initialized away from the private-wealth cap and then simulated under the policy rule implied by the baseline calibration. Notice that the resulting path is not driven by exogenous shocks. It is generated by the endogenous movement of the state variable and the planner's changing willingness to tolerate private-sector expansion. When the private wealth share is low or moderate, the growth mandate makes a relaxed policy attractive: the planner raises private access because doing so delivers higher contemporaneous growth at a relatively small control cost. As this policy increases the private wealth share, however, the economy moves closer to the control cap. The same high-growth policy then becomes costly, and the planner eventually retrenches.

Figure 7: China's GDP path and Model-simulated GDP path



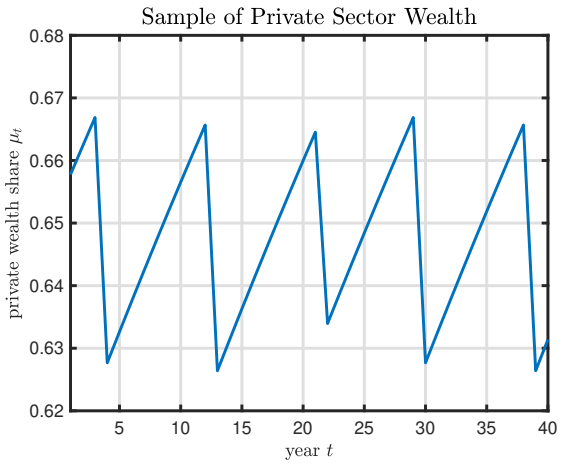
Notes: The left panel plots China's realized GDP path, and the right panel plots the simulated output path under the baseline parameterization. The path is constructed from the model-implied sequence of output growth rates. The path is generated without aggregate shocks.

The model reproduces the qualitative shape of a managed high-growth economy rather than an economy converging smoothly to a balanced-growth path. Growth remains high for a sequence of periods, but it does not converge monotonically to a quiet steady state. Instead, the economy alternates between phases in which the planner pushes growth above its long-run level and phases in which it accepts slower growth in order to restore control. This is the quantitative counterpart of the analytical pendulum result: the same mandate that raises growth in the expansion region also moves the state variable toward the region where retrenchment becomes optimal. A pro-private policy buys growth today by increasing the private wealth share tomorrow. Retrenchment lowers growth today but recreates room for a future expansion.

The private wealth share makes this mechanism visible. In the model, high-growth episodes are

associated with a rising private wealth share, because the relaxed policy allows a larger share of resources to earn the high private return. Retrenchment reverses this process. The swing in the private wealth share is therefore not an auxiliary implication of the model; it is the state variable through which the growth cycle is generated.

Figure 8: Model-Simulated Private Wealth Share



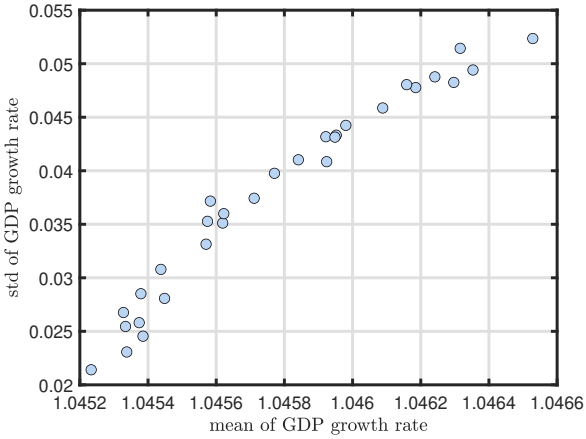
Notes: The figure plots the simulated private wealth share along the same model path as in the previous figure. The recurrent movement of the private wealth share reflects the alternation between growth-supporting relaxation and control-preserving retrenchment.

4.2 The Mean-Volatility Tradeoff

The emergence of the endogenous growth cycle may give rise to a conflict between the mean growth rate and the volatility of the growth rate. If the planner wants to maintain high growth by setting a high growth target, the value of relaxing the private capital restriction when the economy is far from the control cap is high. For a while, the planner can meet the target by allowing more private-sector activity. But this policy also moves the economy closer to the region in which the control motive becomes dominant. Once the state-control concern dominates, the planner has to reverse the earlier relaxation. The higher the target, the larger the accumulated pressure before reversal and, therefore, the sharper the subsequent downturn.

Figure 9 formalizes the mean-volatility tradeoff in tilting the growth target. Each dot in the figure corresponds to one growth target, and the mean and standard deviation of the growth rate are calculated from the simulated endogenous cycles. Raising the mandate increases average growth only by making the economy more cyclical. The mechanism is different from a standard exogenous shock-amplification story. Volatility rises because the policy rule itself becomes more nonlinear. When the target is high, the planner delays retrenchment and keeps the policy stance loose for longer. This increases mean growth in the expansion phase, but it also creates a larger correction once the private wealth share approaches the control region. Thus an aggressive growth target can raise measured growth in normal times while endogenously increasing the probability and severity of a low-growth

Figure 9: Mean versus Volatility



Notes: The figure reports the mean and standard deviation of simulated growth after the burn-in period as the growth mandate is varied around the baseline calibration. Each point is computed from a newly solved policy rule and a simulated path.

episode. This is the model's counterpart to the regional sudden stops in the motivating evidence: when high targets are defended for too long, the eventual correction is not a surprise shock but the delayed resolution of the mandate-control conflict.

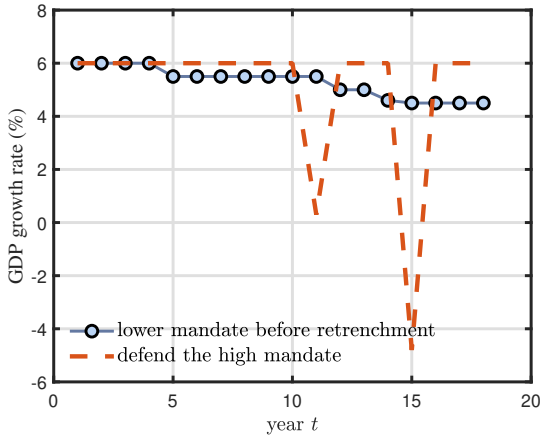
4.3 Soft Landing

The final exercise studies a soft-landing policy and therefore returns to the paper's interpretation of managed deceleration in Figure 2. The baseline high-target economy illustrates the hard-landing force: the planner initially keeps policy loose in order to defend a high mandate, but the resulting rise in the private wealth share eventually makes the high mandate inconsistent with the control objective. If the target is defended until the control motive forces a reversal, the economy experiences a large drop in growth. The soft-landing experiment asks whether the planner can reduce this swing by lowering the growth mandate before retrenchment becomes unavoidable.

Operationally, the experiment solves the model for a menu of growth mandates, from 6.5% down to 4.5%, and compares two simulated paths. The hard-landing path defends the high mandate for an extended period and cuts only after the economy is already close to retrenchment. The soft-landing path lowers the mandate gradually. The point of the exercise is not that a lower target mechanically lowers growth. Rather, a lower target changes the policy rule: it reduces the need to keep private access excessively loose, slows the accumulation of control pressure, and allows the economy to move to a lower-growth path without a dramatic collapse.

The soft-landing result gives a quantitative interpretation of gradual target adjustment in the data. A government that starts from a high target has an incentive to keep policy loose, especially when the private wealth share is not yet close to the cap. Over time, however, the control cost of this policy rises. Lowering the target before the control region is reached relaxes the mandate-control conflict. It allows

Figure 10: Soft Landing through Mandate Adjustment



Notes: The figure compares a path that lowers the growth mandate before retrenchment with a path that defends the high mandate. Growth is reported in percent. The gradual mandate adjustment smooths the transition and avoids the large low-growth episode generated by defending the high target for too long.

the planner to accept a moderate deceleration today in order to avoid a sharp correction tomorrow, exactly the logic suggested by the contrast between national target management and regional sudden stops. In this sense, managed deceleration is not simply a reduced-form preference for smoothness. It is an response to the endogenous state created by past attempts to meet high growth targets.

5. Investment Cycle in Neoclassical Growth Model

The analysis so far has used an AK environment with two ownership sectors. That structure is useful because it maps naturally into the institutional language of state control and private-sector expansion. The economic force, however, is more general than the AK technology. This section shows that a closely related investment cycle arises in a textbook neoclassical growth model once policy is asked to defend a growth mandate with an instrument that also changes the economy's endogenous state. The section proceeds in two steps. First, we lay out the neoclassical environment. Second, we use the quantitative policy rule to show how the model generates alternating episodes of over-investment and under-investment around the standard log-utility saving benchmark.

5.1 A Neoclassical Growth Model with a Growth Mandate

Consider a representative-agent economy with log utility, Cobb-Douglas production, and full depreciation. Time is discrete. At the beginning of period t, the economy inherits aggregate capital k_t. Output is produced according to $y_t = k_t^\alpha$ with $\alpha \in (0,1)$. The planner chooses the saving rate $s_t \in [0,1]$, so consumption is $c_t = (1-s_t)k_t^\alpha$ and next period's capital is$k_{t+1} = s_t k_t^\alpha$. Let the gross policy-relevant growth object be the log change in output,

$$g(k_t, s_t) = \log y_{t+1} - \log y_t = \alpha \log s_t - \alpha(1-\alpha)\log k_t.$$

The planner faces a growth mandate $\bar{g}$ and pays a fixed penalty χ whenever realized growth falls below the mandate. The recursive problem is

$$V(k) = \max_{s \in [0,1]} \left\{ \log\left((1-s)k^{\alpha}\right) - \chi \mathbf{1}\{g(k,s) < \bar{g}\} + \beta V\left(sk^{\alpha}\right) \right\}.$$

The natural benchmark is the standard neoclassical growth model without the mandate. With log utility and full depreciation, the optimal saving rate is constant: $\bar{s} = \alpha\beta$. The steady-state benchmark is therefore governed by the usual Euler logic rather than by target management. The mandate changes the problem because the saving rate now has two roles. It is still the intertemporal allocation instrument that trades off current consumption against future capital, but it is also the direct policy instrument for satisfying the growth target. Conditional on meeting the mandate at capital k, saving must satisfy

$$s \geq \hat{s}(k) \equiv \exp(\bar{g}/\alpha)k^{1-\alpha}.$$

Thus the target-saving schedule $\hat{s}(k)$ is increasing in k. This monotonicity is the key neoclassical force. Because capital has diminishing marginal returns, a given saving rate generates high percentage growth when capital is scarce and low percentage growth when capital is abundant. Defending the same growth mandate therefore requires progressively higher investment as the capital stock rises.

This setup is deliberately close to the canonical model. There is no ownership structure, no AK technology, and no private wealth-share constraint. The only additional ingredient is the target penalty. The question is whether this single friction is enough to produce cyclical investment dynamics. The next subsection shows that it is.

5.2 Investment Cycles

We begin with the induced law of motion for capital. Given the optimal policy rule $s^*(k)$, the capital map is

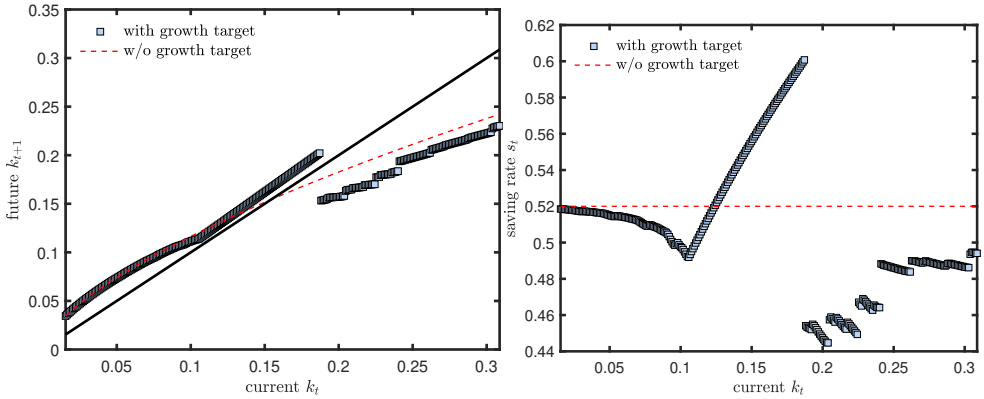
$$K(k) = s^*(k)k^{\alpha}.$$

In the standard log-utility neoclassical model without the mandate, this map is $K^B(k) = \bar{s}k^{\alpha}$. Since $\alpha < 1$, $K^B(k)$ intersects the 45-degree line once at an interior steady state. The red dashed benchmark in the left panel of Figure 11 therefore represents the familiar stable Solow transition.

The growth mandate eliminates this quiet convergence. Around the standard steady state, under a relatively large penalty χ, the planner has an incentive to meet the mandate by raising investment when the benchmark rule would deliver insufficient growth (recall that the growth rate is zero at the standard steady state). That incentive pushes the capital map above the benchmark and, over the relevant region, prevents the economy from settling at the interior fixed point that would prevail absent target management. The resulting dynamics are cyclical rather than monotone.

The right panel of Figure 11 shows the savings policy that generates this capital dynamics. The red dashed line is the constant saving rate $\bar{s}$ in the standard benchmark. We say that periods with $s^*(k) > \bar{s}$

Figure 11: Investment cycle in the neoclassical growth model



Notes: The left panel plots the induced capital law of motion against the 45-degree line. The right panel plots the endogenous saving rate. Red dashed lines report the standard log-utility neoclassical benchmarks: the capital law of motion under $\bar{s} = \alpha\beta$ in the left panel and the constant saving rate $\bar{s} = \alpha\beta$ in the right panel.

are periods of over-investment, while periods with $s^*(k) < \bar{s}$ are periods of under-investment.

The over-investment phase is the direct target-defense margin. When capital is sufficiently high, the benchmark saving rule would imply $g(k, \bar{s}) < \bar{g}$. If the planner chooses to avoid the penalty in that state, it must raise saving at least to $\hat{s}(k)$. Whenever $\hat{s}(k) > \bar{s}$ and the value of avoiding the penalty exceeds the consumption and continuation-value loss from higher saving, the policy rule moves above the neoclassical benchmark.

The same over-investment force, however, changes the state in a way that makes future target defense harder. Since $k_{t+1} = s_t k_t^{\alpha}$, a high saving rate leaves the economy with a larger capital stock tomorrow. Under decreasing returns, that larger inherited capital stock lowers the growth rate associated with any fixed saving rate and raises the future target-saving threshold $\hat{s}(k_{t+1})$. Hence over-investment is not simply a one-period distortion. It accumulates the state that eventually makes continued target defense increasingly costly. This is the neoclassical analogue of the AK model's control pressure: in the baseline model, pro-private policy raises the future private wealth share; here, high investment raises the future capital stock and thereby worsens the growth arithmetic.

The under-investment phase is interesting and contains two regions: there will be under-investment when the capital stock is very high or when it is very low. Under-investment in times of high capital stock is intuitive. In those states, meeting the growth mandate requires very high saving rate, if not impossible. This greatly distorts the intertemporal consumption-saving substitution, and the loss exceeds the penalty of missing the growth mandate. Therefore, the planner would miss the target. However, instead of following the benchmar saving rate $\bar{s}$, the planner would set a even lower saving rate $s^*(k_t) < \bar{s}$. This lowers the capital stock k_{t+1} in the next period and corrects the large intertemporal distortion.

Under-investment in times of low capital stock reflects a forward-looking concern. Suppose the continuation problem contains a region in which the planner expects to defend the mandate. Entering

that region with a higher capital stock raises the future required saving rate $\hat{s}(k)$ and magnifies the distortion needed to meet the target. A forward-looking planner can reduce this future distortion by saving less today. Lower saving reduces k_{t+1}, making next period's percentage growth target easier to meet and reducing the future deviation from $\bar{s}$. Meanwhile, because the current capital stock is low, doing this can still meet the growth mandate today. Thus under-investment in those states is not the opposite of target management. It is part of target management: the planner deliberately holds down the state variable in order to make future target compliance less costly.

The discreteness of the mandate penalty helps explain the irregular shape of the policy rule and the law of motion. At each state, the planner compares two regimes: pay the fixed cost of missing the target and choose a relatively undistorted saving rate, or meet the target and accept the consumption and intertemporal losses from a distorted saving rate. The boundary between these regimes creates kinks in the continuation value. Because today's saving determines tomorrow's location relative to those kinks, the future meet-or-miss decision is propagated backward into current investment. This is why the saving rule in Figure 11 is not a smooth monotone tilt above the neoclassical benchmark. It contains regions in which the planner deliberately keeps capital lower, followed by regions in which it invests aggressively to maintain the target, and then discrete adjustments when the target becomes too costly to defend.

The neoclassical extension therefore clarifies the robustness of the paper's mechanism. What is essential is not the AK assumption, nor the particular interpretation of the state variable as a private wealth share. What is essential is a growth mandate attached to a policy instrument that changes the state on which future mandate feasibility depends. In the baseline model, the state is the private wealth share and the relevant future constraint is political control. In the neoclassical model, the state is aggregate capital and the relevant force is diminishing returns. In both cases, defending growth today changes tomorrow's environment in a direction that eventually makes continued defense costly.

6. Conclusion

This paper develops a unified framework for understanding why economic policy often moves in cycles between market liberalization and state control. The central mechanism is a structural tension between growth and control: policies that relax constraints on more productive private activity help satisfy growth mandates, but they also increase the private wealth share and eventually threaten the planner's control objective. Once this constraint becomes sufficiently tight, the planner optimally reverses course, generating retrenchment and lower growth before the pressure for expansion reemerges. The model shows that such policy volatility need not arise from exogenous shocks or political turnover; it can be an endogenous consequence of governing under binding growth targets. The empirical evidence from China's post-reform experience, together with supplementary evidence from U.S. financial cycles, suggests that this trade-off is relevant across institutional settings. The analysis also highlights an important policy implication: aggressive growth targets may raise aver-

age growth only at the cost of greater volatility, whereas a gradual reduction in growth mandates can facilitate a smoother transition and reduce the risk of sharp corrections.

References

Alesina, A. (1987). Macroeconomic policy in a two-party system as a repeated game. *Quarterly Journal of Economics 102*(3), 651–678.

Beaudry, P., D. Galizia, and F. Portier (2020). Putting the cycle back into business cycle analysis. *American Economic Review 110*(1), 1–47.

Benhabib, J. and R. E. A. Farmer (1994). Indeterminacy and increasing returns. *Journal of Economic Theory 63*(1), 19–41.

Benhabib, J. and K. Nishimura (1979). The hopf bifurcation and the existence and stability of closed orbits in multisector models of optimal economic growth. *Journal of Economic Theory 21*(3), 421–444.

Boldrin, M. and A. Rustichini (1994). Indeterminacy of equilibria in models with infinitely-lived agents and external effects. *Econometrica 62*(2), 323–342.

Brandt, L. and X. Zhu (2000). Redistribution in a decentralized economy: Growth and inflation in china under reform. *Journal of Political Economy 108*(2), 422–439.

Cass, D. and K. Shell (1983). Do sunspots matter? *Journal of Political Economy 91*(2), 193–227.

Cheremukhin, A., M. Golosov, S. Guriev, and A. Tsyvinski (2024, April). The political development cycle: The right and the left in people's republic of china from 1953. *American Economic Review 114*(4), 1107–1139.

Christiano, L. J. and S. G. Harrison (1999). Chaos, sunspots and automatic stabilizers. *Journal of Monetary Economics 44*(1), 3–31.

Hibbs, D. A. (1977). Political parties and macroeconomic policy. *American Political Science Review 71*(4), 1467–1487.

Lin, J. Y., F. Cai, and Z. Li (1996). *The China Miracle: Development Strategy and Economic Reform*. Hong Kong: Chinese University Press.

Naughton, B. (1996). *Growing Out of the Plan: Chinese Economic Reform, 1978–1993*. Cambridge: Cambridge University Press.

Nordhaus, W. D. (1975). The political business cycle. *Review of Economic Studies 42*(2), 169–190.

Persson, T. and G. Tabellini (1990). *Macroeconomic Policy, Credibility and Politics*. Chur, Switzerland: Harwood Academic Publishers.

Rogoff, K. (1990). Equilibrium political budget cycles. *American Economic Review 80*(1), 21–36.

Wang, Y. (2021). The politico-economic dynamics of china's growth. *Journal of the European Economic Association 19*(6), 3324–3369.

Woodford, M. (1991). Self-fulfilling expectations and fluctuations in aggregate demand. *American Economic Review 81*(4), 921–936.

Xu, Z., J. Yu, and Z. Zhang (2025). A political model of china's macroeconomy. *Available at SSRN 5328567*.

Appendix

Proof of Lemma 3.1

Proof. For Lemma 3.1. Fix a household $j \in \{p, s\}$ and take the sequence of returns $\{R_{j,t}\}_{t\geq 0}$ as given. The household's budget constraint is

$$c_{j,t} + a_{j,t+1} = R_{j,t} a_{j,t}.$$

With log utility, the Euler equation for an interior optimum is

$$\frac{1}{c_{j,t}} = \beta_j R_{j,t+1} \frac{1}{c_{j,t+1}}.$$

Consider the policy

$$c_{j,t} = (1-\beta_j) R_{j,t} a_{j,t}, \qquad a_{j,t+1} = \beta_j R_{j,t} a_{j,t}.$$

It satisfies the budget constraint. Moreover, because $c_{j,t+1} = (1-\beta_j) R_{j,t+1} a_{j,t+1}$,

$$\beta_j R_{j,t+1} \frac{1}{c_{j,t+1}} = \frac{\beta_j R_{j,t+1}}{(1-\beta_j) R_{j,t+1} a_{j,t+1}} = \frac{\beta_j}{(1-\beta_j) a_{j,t+1}} = \frac{1}{(1-\beta_j) R_{j,t} a_{j,t}} = \frac{1}{c_{j,t}}.$$

Thus the Euler equation holds, and the usual transversality condition holds under the log-saving rule. Hence each household saves the constant fraction β_j of its end-of-period wealth. For the state household, $R_s = z_s$, so

$$a_{s,t+1} = \beta_s R_s a_{s,t}.$$

For the private household, the effective return on beginning-of-period wealth is

$$R_p(\lambda_t) = \lambda_t z_p + (1-\lambda_t) z_s,$$

because the fraction λ_t is operated with the private technology and the remaining fraction $1-\lambda_t$ earns the state return. Therefore

$$a_{p,t+1} = \beta_p R_p(\lambda_t) a_{p,t}.$$

□

Proof of Lemma 3.2

Proof. For Lemma 3.2. Let $K_t = a_{p,t} + a_{s,t}$ and $\mu_t = a_{p,t}/K_t$. By Lemma 3.1,

$$a_{p,t+1} = \beta_p R_p(\lambda_t) a_{p,t}, \qquad a_{s,t+1} = \beta_s R_s a_{s,t}.$$

Therefore

$$\begin{aligned}\mu_{t+1} &= \frac{a_{p,t+1}}{a_{p,t+1} + a_{s,t+1}} \\ &= \frac{\beta_p R_p(\lambda_t) a_{p,t}}{\beta_p R_p(\lambda_t) a_{p,t} + \beta_s R_s a_{s,t}} \\ &= \frac{\mu_t \beta_p R_p(\lambda_t)}{(1-\mu_t)\beta_s R_s + \mu_t \beta_p R_p(\lambda_t)} = \mathcal{M}(\mu_t, \lambda_t).\end{aligned}$$

Next, current output is

$$\begin{aligned}Y_t &= z_p \lambda_t a_{p,t} + z_s \left[a_{s,t} + (1-\lambda_t) a_{p,t}\right] \\ &= \left[z_s + (z_p - z_s)\lambda_t \mu_t\right] K_t.\end{aligned}$$

To express Y_{t-1} in terms of period-t wealth, use the saving rules:

$$a_{p,t} = \beta_p \,\text{income}_{p,t-1}, \qquad a_{s,t} = \beta_s \,\text{income}_{s,t-1}.$$

Since total output in period $t-1$ is total pre-saving income,

$$\begin{aligned}Y_{t-1} &= \frac{a_{p,t}}{\beta_p} + \frac{a_{s,t}}{\beta_s} \\ &= K_t \left(\frac{\mu_t}{\beta_p} + \frac{1-\mu_t}{\beta_s}\right) \\ &= K_t \left(\frac{1}{\beta_s} + \frac{\beta_s - \beta_p}{\beta_s \beta_p} \mu_t\right).\end{aligned}$$

Hence

$$g_t = \frac{Y_t}{Y_{t-1}} = \frac{z_s + (z_p - z_s)\lambda_t \mu_t}{\frac{1}{\beta_s} + \frac{\beta_s - \beta_p}{\beta_s \beta_p} \mu_t} = \mathcal{G}(\mu_t, \lambda_t).$$

It remains to verify the monotonicity statements. Write$A(\lambda) = \beta_p R_p(\lambda)$ and $B = \beta_s R_s$. Then

$$\mathcal{M}(\mu, \lambda) = \frac{\mu A(\lambda)}{(1-\mu)B + \mu A(\lambda)}.$$

Since $A(\lambda) > 0$, $B > 0$, and $A'(\lambda) = \beta_p(z_p - z_s) > 0$,

$$\frac{\partial \mathcal{M}}{\partial \mu} = \frac{A(\lambda) B}{\left[(1-\mu)B + \mu A(\lambda)\right]^2} > 0, \qquad \frac{\partial \mathcal{M}}{\partial \lambda} = \frac{\mu(1-\mu) B A'(\lambda)}{\left[(1-\mu)B + \mu A(\lambda)\right]^2} > 0$$

for $\mu \in (0,1)$. Finally,

$$\frac{\partial \mathcal{G}}{\partial \lambda} = \frac{(z_p - z_s)\mu}{\frac{1}{\beta_s} + \frac{\beta_s - \beta_p}{\beta_s \beta_p} \mu} > 0$$

for $\mu > 0$, because $z_p > z_s$ and the denominator is positive. □