Who Bears the Cycle? Time-Consistent Fiscal Policy and the Financing Assignment in Emerging Markets*

Nickson J. Cabote†

August 14, 2026

Draft — do not cite

Abstract

Fiscal rules are often recommended to limit procyclical spending in emerging and developing economies. This paper looks at these rules using a calibrated small open-economy model. In the model, the government cannot commit in advance, so it chooses public spending each period and faces shocks to productivity, world interest rates, and household preferences. Results show that, despite the absence of political or default risks, procyclical spending remains the Markov-perfect best response. With public debt held constant, cyclical rules only shift the adjustment between spending and taxes rather than remove it. In this model, the main cost of discretion is the overall level of public purchases, not how it changes during the cycle. Sustained restraint would improve welfare, but a government making decisions year by year will not maintain it. Cyclical rules mostly shift which periods bear the cost. Rules focused on spending cause losses during times of tight fiscal limits, while rules focused on financing shift losses to high-productivity periods. In all the cases we study, the effect on the overall level is greater than the effect of shifting costs between periods. ***JEL classification:*** *E62; E63; F41; H30; H63.*

Keywords: *time-consistent fiscal policy; financing assignment; emerging markets; small open economy; Markov-perfect equilibrium; consumption-equivalent welfare.*

*I am grateful to Professor Jinhui Bai for invaluable guidance throughout this project, and to Professor Salvador Ortigueira and Professor Jia Yan for valuable comments at the Final Research Defense (March 2026). Earlier versions were presented at the Washington State University School of Economic Sciences Friday Seminar (Pullman, WA, November 2025) and at the Western Economic Association International (WEAI) International Conference (Bangkok, Thailand, April 2026). I thank seminar and conference participants for helpful comments. The views in this paper are those of the author and do not necessarily represent those of the Bangko Sentral ng Pilipinas. All remaining errors are my own.

†Researcher, Bangko Sentral ng Pilipinas. Email: `CaboteNJ@bsp.gov.ph`.

Acknowledgements and Disclosures

Who Bears the Cycle? Time-Consistent Fiscal Policy and the Financing Assignment in Emerging Markets

Nickson J. Cabote · August 14, 2026

I am grateful to Professor Jinhui Bai for invaluable guidance throughout this project, and to Professor Salvador Ortigueira and Professor Jia Yan for valuable comments at the Final Research Defense (March 2026). Earlier versions were presented at the Washington State University School of Economic Sciences Friday Seminar (Pullman, WA, November 2025) and at the Western Economic Association International (WEAI) International Conference (Bangkok, Thailand, April 2026). I thank seminar and conference participants for helpful comments.

Research support. This paper is based on a chapter of the author's doctoral dissertation at Washington State University. The research received no external grant funding. Institutional support for the dissertation was provided by the School of Economic Sciences at Washington State University.

Financial relationships and potential conflicts of interest. The author is employed as a Researcher at the Bangko Sentral ng Pilipinas (BSP) Research Academy. That employment is a material financial relationship under the NBER research dissemination disclosure policy. The author has also previously held macroeconomic policy positions at the BSP (Center for Monetary and Financial Policy) and at the Department of Economy, Planning, and Development of the Philippines. The Bank did not provide dedicated project funding, did not have the right to review this paper prior to circulation, and is not responsible for the analysis or the conclusions. The views in this paper are those of the author and do not necessarily represent those of the Bangko Sentral ng Pilipinas or Washington State University.

The author declares that he has no other relevant or material financial interests that relate to the research described in this paper, including consulting, ownership, or paid advisory relationships active in the last three years. All remaining errors are my own.

1 Introduction

Fiscal policy in emerging markets is usually procyclical, meaning government spending rises during economic booms and falls during downturns. Researchers have seen this pattern in many regions and with different methods for over thirty years (Gavin and Perotti, 1997; Kaminsky et al., 2005; Ilzetzki and Végh, 2008; Frankel et al., 2013), and it is often viewed as a policy concern. Most explanations focus on political-economy problems (Tornell and Lane, 1999; Alesina et al., 2008) or challenges with external finance and the risk of sovereign default (Cuadra et al., 2010; Bianchi et al., 2023). These factors create obstacles that are not present in the standard time-consistent model. In that model, households and firms make the best choices they can, and the government aims to maximize welfare but cannot promise future policies.

We take this standard benchmark to the data. Our approach uses a calibrated, time-consistent small open economy model for emerging markets, where a Markov-perfect fiscal authority decides on government spending each period in response to shocks to productivity, the world real interest rate, and preferences. This model does not include political-economy factors, sovereign default risk, or a debt-dependent interest rate premium. The Markov-perfect equilibrium in our model leads to procyclical government spending as the optimal response. Spending and output move together, and relative volatility is close to the data ($\sigma(\tilde{G})/\sigma(\tilde{Y}) = 1.94$ against 2.20); Section 6.2 reports the full comparison. The only fiscal limits in the model are limited commitment, constant debt, and a ceiling on next-period purchases relative to next-period output. This approach adds to, rather than replacing, the political-economy and default explanations found in other research.

Using this microfounded benchmark, we evaluate the welfare cost of cyclical fiscal-rule reforms. The standard counterfactual the literature uses smoothes or makes counter-cyclical the path of government purchases and reads the welfare difference relative to the procyclical equilibrium. This counterfactual is incomplete when the level of public debt is held stable and the labor tax that finances government purchases adjusts to satisfy the period budget. When government purchases are smoothed over the cycle, the adjustment shifts to the financing side. This makes the correlation between the implementing tax and output much more negative than in the benchmark (-0.49 compared to -0.33). The welfare comparison is therefore not between two cyclical rules on government purchases but between two joint packages of purchases and the implementing tax. The appropriate policy variable in welfare comparisons of cyclical fiscal-rule reforms is the assignment of cyclical adjustment between the spending and financing sides, not the cyclical stance of either side in isolation. Lemma 1 explains this duality in analytical terms. The next Remark builds on this, showing that the way cyclical adjustment is assigned does not depend on using constant-debt closure in our experiment. This result holds for any cyclical debt-issuance rule, except in the unlikely case of a perfectly state-contingent elasticity that absorbs all cyclical financing changes, which is not realistic in emerging markets due to market discipline and how rules are designed.

When a counter-cyclical rule is used for purchases, as in the typical emerging-market policy scenario, welfare falls by about 1.2% in consumption-equivalent terms on the headline discretization. The corresponding output-indexed tax rule also lowers welfare, by about 0.5%, as do the two

constant-level rules. These four ergodic cells are not separated from one another at the precision the computations support (Section 6.3). What the assignment determines, and what those computations do identify, is which states bear the cost: the two sides are close to mirror images. The spending-side loss is concentrated in the low-productivity, high-world-rate quadrant where fiscal feasibility tightens, at about -12.6% of consumption (cross-seed sd about 0.6pp) against near-neutral values in the common low-rate states. The financing-side rule has the opposite incidence: it gains in those common low-rate states and loses in high-productivity ones. Splitting each CEV into its private- and public-consumption channels shows the mechanism (see Section 6.3). Reversing the purchases cycle removes public-consumption value while returning almost nothing to the household. The output-indexed tax rule concedes less public value and returns a small amount. The constant-purchases cell is dominated by its level rather than by its constancy. Appendix I reports the grids and the identification limits.

Setting the cyclical comparison beside a level comparison shows that the cycle is the smaller part of what discretion costs. The discretionary government provides public goods at 31% of output. The welfare-maximizing spending share, computed period by period from the equilibrium's own absorption, puts purchases at 17%. Permanent proportional restraint—scaling the equilibrium's state-contingent policy down by a constant factor, which preserves its cyclical shape exactly and moves only its level—traces a welfare gain that peaks at $+13.9\%$ of consumption at a mean G/Y of 0.172, within four thousandths of that independently computed share. That is an order of magnitude larger than any of the cyclical cells, and it is estimated far more precisely, with a cross-seed standard deviation of 0.34.

A government choosing one period at a time will not move there, and the reason is the object this paper is about. Restraint pays only if it is sustained. At the same scale factor, a single period of lower purchases is worth -0.47% by the government's own objective, about one thirtieth of what sustaining it pays, because next period's government inherits the state and reverts to the equilibrium policy. The deviation buys one period of a different split and nothing afterward. The gap between what sustained restraint is worth and what unilateral restraint is worth is the value of commitment, and it is measured here in the same consumption-equivalent units as the rule comparisons. It also has a numerical counterpart, which is the same fact viewed from the solver's side: the government's objective is nearly flat in its choice of purchases, exactly as it must be when a one-period deviation is nearly worthless. That flatness is why we quote the cyclical cells as a band rather than as point estimates.

Output is determined by productivity and the world interest rate through the static capital-demand condition, with labor supply fixed, so Y_t is invariant to the fiscal rule. The welfare comparison therefore isolates the assignment channel: it does not net a stabilization benefit against an allocation cost. A negative CEV is then an upper bound on the net cost of the same rule in an environment with a demand channel, and a direct measure of the financing-assignment cost by itself.

Here is how the mechanism works. When productivity is high or the world interest rate is

low, the tax base grows and the marginal utility of private consumption is low, so the real cost of financing an extra unit of purchases is also low. As a result, the government increases spending, but only up to the ceiling $G_{t+1} \leq \bar{g}\, Y_{t+1}$. When conditions worsen, people spend less, and the value they place on extra consumption goes up. This means that the same level of government spending costs households more in terms of welfare, and taxes also have to increase. So, the government reacts by raising taxes and reducing spending. Procyclical government spending happens naturally because of limited commitment, the changing welfare cost of fiscal financing, and that ceiling.

World-rate fragmentation and the assignment problem. Cyclical adjustment has become a key topic in discussions about how to design fiscal rules for emerging markets facing global financial fragmentation. In small open economies, global financial conditions affect domestic fiscal feasibility mainly through the world real interest rate. When geopolitical fragmentation increases, funding costs become more varied, or the world interest rate rises and becomes more volatile, the variability of external financing grows just as fiscal feasibility becomes more limited. When external financing becomes more unpredictable, it makes sense to focus on cyclical adjustment as a policy tool. The cost of adjusting spending is highest in times of low productivity and high world interest rates; these are also the times when fragmentation makes fiscal feasibility more difficult. In these situations, a reform that stabilizes public spending over the cycle is essentially the same as a reform that makes labor taxes less stable, especially when this is most expensive.

We use annual data from the Philippines from 2000 to 2019 to set up our quantitative analysis, rather than to provide a detailed country case study. According to standard measures of openness, external balance-sheet exposure, and institutional quality, the Philippines is about average among emerging markets (World Bank, 2025a; Lane and Milesi-Ferretti, 2018; World Bank, 2025b). Its fiscal data also show the same patterns of procyclicality and volatility seen in other countries (Kaminsky et al., 2005; Talvi and Végh, 2005; Frankel et al., 2013; Vegh and Vuletin, 2015). The period from 2000 to 2019 covers the aftermath of the Asian financial crisis, the global financial crisis, and the post-2010 infrastructure expansion. During these years, limited fiscal space and external financing often made it difficult for the government to increase spending (Cevik, 2019; Alberola et al., 2021; International Monetary Fund Independent Evaluation Office, 2002).

Since our welfare comparison relies on numerical equilibria, we state plainly which results the computations support outright and which we report with their uncertainty. Three findings are established at the resolution we can compute. First, discretion over-provides public goods in levels, and closing that gap by permanent proportional restraint is worth +13.9% of consumption with a cross-seed standard deviation of 0.34; two computations that share no code, a welfare-maximizing spending share and a dynamic re-solve with the household responding, locate the preferred share at G/Y of 0.168 and 0.172. Second, sustained restraint is worth about thirty times a single period of the same restraint by the government's own objective—the value of commitment—which identifies the level gap as a commitment problem rather than a solver artifact. Third, the two sides of the budget have opposite state-conditional incidence: the spending-side rule costs −12.6% in the low-

productivity, high-world-rate quadrant where fiscal feasibility tightens, while the financing-side rule costs −13.0% in high-productivity states and gains +4.3% in the low-rate states that make up most of the ergodic distribution. All three are large relative to their cross-seed standard deviations. The first two are also large relative to the discretization uncertainty documented in Appendix C; the third is identified within the headline economy, where the state-conditional cells were computed.

One further quantity is quoted as a band. The ergodic magnitudes of the four rule cells inherit the flatness of the government's objective—the median on-path separation between the best and second-best purchases node is worth 0.10% of consumption on the headline grid (0.02% after the node re-allocation)—so we give them to the nearest half point and rest no argument on differences smaller than a point, and comparisons involving the constant-purchases rule carry the additional level qualification of Section 6.3. Appendix C quantifies the flatness and Appendix I traces its consequences. Both appendices also show what the first two findings are invariant to, including the asset normalization: dropping the household's claims on capital and public debt ($A_{\text{ref}} = 0$) moves the reference labor tax to 0.243 and the reference net foreign asset position to −268% of output while leaving the level ranking and the value of commitment unchanged.

Section 2 relates our work to research on fiscal cyclicality and emerging-market business cycles. Section 3 documents the cyclical pattern of Philippine government purchases. Section 4 presents the model, the government's problem, and the analytical duality between the cyclical stances of purchases and the financing tax. Section 5 describes the calibration and the global solution method. Section 6 reports the fitted policy rules, simulated business-cycle properties, the model fit to emerging-market data, and the welfare comparison. Appendix D records the fixed-point iteration. Appendix I reports robustness, and Section 7 concludes.

2 Related Literature

This paper brings together three areas of research: empirical studies on fiscal procyclicality in emerging markets, structural models of how fiscal policy is determined, and the optimal-policy approaches of Ramsey and Markov-perfect taxation.

The empirical literature establishes the regularity we take to the model. Fiscal policy in emerging markets is procyclical across regions and identification strategies, with positive comovement between expenditures and output (Gavin and Perotti, 1997; Kaminsky et al., 2005; Ilzetzki and Végh, 2008; Frankel et al., 2013).[1] Previous research has looked at how procyclical fiscal policy can worsen economic cycles and slow growth (Aghion et al., 2007; Fatás and Mihov, 2013; Loayza and Raddatz, 2007). About a third of emerging markets have shifted to acyclical or countercyclical policies after adopting fiscal rules (Frankel, 2011; Jalles et al., 2024). In our study, we focus on government purchases. We use spending aggregates and do not include transfers in our measure of G (Lane, 2003; Marioli et al., 2024). Researchers usually mention two main reasons for this

[1] The robustness of the pattern to instrumental-variable, forecast-error, and panel approaches is documented by Talvi and Végh (2005); Ilzetzki et al. (2013); Jaimovich and Panizza (2007); Alesina et al. (2008).

pattern. First, when it is harder to get external financing in downturns, governments have to cut spending at the same time that stimulus is needed (Gavin and Perotti, 1997; Calvo et al., 2004; Bianchi et al., 2023).

Second, political-economy frictions like common-pool problems and rent-seeking lead to spending that grows faster than the economy during booms (Tornell and Lane, 1999; Alesina et al., 2008; Woo, 2009).[2] Research on sovereign default (Aguiar and Gopinath, 2007; Arellano, 2008; Cuadra et al., 2010; Bianchi et al., 2023) points to another reason for procyclical fiscal policy in emerging markets: during downturns, higher sovereign spreads make borrowing too expensive, so governments are forced to cut spending. Our identified channel is different from this. Even if there is no risk of default, and with limited commitment and residual labor-tax financing, the time-consistent policy for government purchases still ends up being procyclical, and the financing-assignment cost we discuss in Section 6.3 appears. Both mechanisms can happen at the same time. If we added default risk to our model, government spending would become even more procyclical in downturns because of higher spreads, and the welfare cost of spending rules would rise due to greater deadweight loss from default-related spreads. Still, the overall direction of the financing-assignment asymmetry would not change.

On the modeling side, the most relevant earlier work comes from Bachmann and Bai (Bachmann and Bai, 2013b,a), who study a welfare-maximizing, time-consistent authority that chooses public consumption while dealing with implementation frictions. They find that, to match the volatility and persistence of public spending seen in the data, it is necessary to include adjustment costs and taste shocks. However, their studies do not offer practical counterfactuals for how government spending should respond to the business cycle when labor taxes are used to balance the budget under a simplified fiscal rule. In this paper, we use the same policy framework but ask a different welfare question: how do different rules for the cyclicality of G compare when the way G is financed through τ_ℓ is made clear in an open economy? The main finding from earlier work is that a government without frictions produces too little volatility and persistence to match real-world data unless there are adjustment costs and preference shocks. This result also applies to emerging-market open economies. Our model includes these elements and places them in a small open economy with incomplete international asset markets (Mendoza, 1991; Schmitt-Grohé and Uribe, 2003), as well as shocks to productivity, preferences, and the world real interest rate (Neumeyer and Perri, 2005; Uribe and Yue, 2006; Aguiar and Gopinath, 2007). When compared to Ramsey benchmarks with complete markets and full commitment (Barro, 1979; Lucas and Stokey, 1983), the Markov-perfect approach shows how limited commitment and features of emerging markets affect the cyclicality of G. This study also contributes to the broader literature on incomplete commitment (Maskin and Tirole, 2001; Klein et al., 2008; Debortoli and Nunes, 2010).

The financing side of the welfare comparison—which part of the budget adjusts cyclically under residual financing—has been analyzed in the optimal-taxation literature without state-contingent

[2]The interaction between the two channels and the role of rule design (rather than rule adoption) in mitigating the cyclical bias is discussed in Riascos and Vegh (2003); Cuadra et al. (2010); Alt and Lassen (2006); Schaechter et al. (2012); Bova et al. (2014); Guerguil et al. (2017); Mendoza and Oviedo (2006).

debt. Aiyagari et al. (2002) describe optimal labor taxation when the planner can issue only risk-free debt; in their model, the optimal tax has a highly persistent component because the financing residual cannot be fully smoothed. Werning (2007) extends this to a redistributive Ramsey planner and describes tax-smoothing conditions on the financing side. Pouzo and Presno (2022) solve an optimal-taxation problem with endogenous default, where the credit limit interacts with the tax-smoothing margin. These studies point out that the financing mix is important for optimal-taxation problems in incomplete markets, but they focus on Ramsey commitment and do not consider how spending changes with the economic cycle. In this study, which looks at emerging markets with time-consistent policies, the way financing is assigned reflects the welfare cost of one cyclical-rule reform compared to the other side of the budget. Most cross-country research on fiscal cyclicality in emerging markets does not clearly explain how financing is assigned. Instead, these studies usually compare a procyclical rule with a countercyclical rule without saying how the alternative is funded. One study that considers both aspects is Cuadra et al. (2010).

The welfare comparison focuses on two main points. First, the allocations can be carried out and remain consistent over time under the same constant-debt rule and a common long-run average of G/Y in the cyclical-stress experiments. Second, the financing path is part of the allocation. Because labor taxes are determined as a residual under the constant-debt rule, the model itself sets the cyclicality of τ_ℓ, so it cannot be chosen separately. The Philippines is the calibration target because it is close to the middle of the emerging-market range for income, openness, and institutional quality. Fiscal data from 2000 to 2019 for the Philippines show the same procyclicality and volatility found in cross-country studies.

3 Fiscal Cyclicality in Emerging Markets: Evidence

This section describes the cyclical pattern of government purchases that the model aims to match. We use annual Philippine public finance data from 2000 to 2019. The key statistics fall within published ranges for emerging markets and focus on three main aspects: the comovement, persistence, and volatility of G. In this context, the Philippines is a calibration target, not a case study. Kaminsky et al. (2005); Talvi and Végh (2005); Frankel et al. (2013); Vegh and Vuletin (2015) show that these statistics are well within the interquartile bands for emerging markets.

Definitions and notation. We define G (government purchases) as government consumption expenditures plus gross public investment, but not transfers, subsidies, or interest payments. This total includes collective services like defense and public safety, individual services provided by the government such as public education and health, and infrastructure. This approach matches the NIPA category called "government consumption expenditures and gross investment" (U.S. Bureau of Economic Analysis, 2024), where social benefits and interest are not part of G. Y stands for real GDP. Procyclicality refers to the correlation between the cyclical parts of Y and G at the same time. If $\text{corr}(\tilde{Y}, \tilde{G}) > 0$, spending is procyclical. If the correlation is about zero, spending is

acyclical. If it is less than zero, spending is countercyclical. Here, $\tilde{X}$ means the HP-detrended log series. It is important to keep purchases and transfers separate.

Purchases are procyclical in advanced economies as well (Lane, 2003), but they are much more volatile in emerging markets (Ilzetzki and Végh, 2008; Marioli et al., 2024). In our model, since public services affect utility, the way G moves with the cycle determines how much risk-sharing the fiscal authority can provide. If G is procyclical, public goods shrink exactly when people need them most, which limits the extra insurance G can offer (Riascos and Vegh, 2003; Pallage and Robe, 2003).

Data construction. We construct annual fiscal aggregates using the Department of Budget and Management's Budget of Expenditures and Sources of Financing (BESF; various years). These are mapped to Government Finance Statistics Manual (GFSM 2014) categories, following the approach in the IMF Fiscal Transparency Evaluation for the Philippines (Department of Budget and Management, 2025; International Monetary Fund, 2014; Mueller et al., 2015).[3] For output Y, we use annual real GDP from the Philippine Statistics Authority (Philippine Statistics Authority, 2020). Before detrending, we use log levels and define $\tilde{X}$ as the cyclical component from an HP filter with $\lambda = 100$ for the years 2000 to 2019.

Moments and representativeness. Table 1 presents the volatility, persistence, and comovement of key fiscal aggregates. Government purchases G have a standard deviation of 6.6%, an AR(1) persistence of 0.36, and a correlation with output of 0.34.[4] Economic services, including public investment, are about twice as volatile as G and are only weakly procyclical. Social services have volatility similar to G but are more strongly procyclical. The values for G are within the interquartile ranges found in cross-country studies of emerging-market fiscal cyclicality (Kaminsky et al., 2005; Talvi and Végh, 2005; Frankel et al., 2013; Vegh and Vuletin, 2015). This supports using the Philippines as a calibration target instead of treating it as a special case. Net lending and interest payments are excluded so the series match the model's G.

Robustness to detrending. For both HP ($\lambda = 6.25$ and 100) and Hamilton filters, G is volatile, moderately persistent, and procyclical. The Baxter–King filter shows weaker contemporaneous comovement in the short annual sample, which is due to the well-known endpoint loss of band-pass methods.

[3] Fiscal aggregates are based on BESF obligations on a cash basis, deflated to constant 2010 prices to match the PSA real-GDP measure for Y.

[4] Philippine-specific evidence on the cyclical fiscal *stance* is more mixed than this comovement alone suggests. Cevik (2019) finds a procyclical discretionary stance through 2016 using the cyclically adjusted primary balance, while Debuque-Gonzales et al. (2026) estimate a Bohn-type fiscal reaction function in which the output-gap response is imprecisely estimated and changes sign across specifications. That evidence measures a different object from the moment used here: the response of a cyclically adjusted balance to the output gap, rather than the comovement of government purchases with output. The model is calibrated to the latter, which is also the object of the cross-country emerging-market evidence cited above.

Table 1: Philippines (annual), 2000–2019: Core fiscal aggregates vs. GDP

	G	$G^{\text{non-def}}$	Y	ES	SS
Standard deviation (%)	6.6	6.8	3.0	12.6	6.2
Relative volatility $\sigma(X)/\sigma(Y)$	2.20	2.27	1.00	4.20	2.07
AR(1) persistence ρ	0.36	0.31	0.67	0.37	0.26
Correlation with output	0.34	0.33	1.00	0.21	0.47

Acronyms: Y = output (GDP); ES = economic services (public investment); G = government purchases; $G^{\text{non-def}}$ = government purchases excluding defense; SS = social services.
Notes: Annual log deviations from trend, HP filter $\lambda = 100$, sample 2000–2019 ($T = 20$). Standard deviations are reported in percent. Correlations are between cyclical components. AR(1) persistence ρ is from $\tilde{X}_t = \rho\,\tilde{X}_{t-1} + \varepsilon_t$. With $T = 20$, sampling error is large: analytic 90% intervals are approximately $[-0.04, +0.64]$ for $\text{corr}(\tilde{G}, \tilde{Y})$ and $[0.02, 0.70]$ for the AR(1) of $\tilde{G}$ (Fisher / OLS approximations; see the replication package). Absolute volatility intervals are tighter, about $[4.8, 8.4]$ for $\sigma(\tilde{G})$ and $[2.2, 3.8]$ for $\sigma(\tilde{Y})$. Net lending and interest payments are omitted so the series match model G. Sources: DBM *Budget of Expenditures and Sources of Financing*; PSA *National Accounts*.

Regularities that discipline the quantitative model. Two main properties of G are important for emerging markets and help us set our calibration. First, G usually moves in the same direction as output each year and tends to stay that way over time. This matches what we see with multi-year budget cycles and delays in putting plans into action (Leeper et al., 2010; International Monetary Fund, 2019). Second, G is more volatile than Y, with a volatility ratio above one. Table 1 shows that economic services, like public investment, are about twice as volatile as total purchases and often shrink more during downturns unless specific rules protect them. Social services are about as volatile as G but are even more procyclical (Frankel et al., 2013; Ardanaz et al., 2020).

These patterns help set the model's empirical goals, like tracking how G moves and persists over time and how much of its variance is explained by macro fundamentals. They also support using the standard shock structure with productivity, world interest rates, and preferences, as seen in emerging-market SOE research (Aguiar and Gopinath, 2007; Neumeyer and Perri, 2005; Uribe and Yue, 2006).

4 Model

We study a small open economy with flexible prices, a representative household, competitive firms, and a fiscal authority that chooses government spending but cannot fully commit to future plans. Our model stands out from the standard real business-cycle small open economy in three main ways. First, the government sets G each period as its best response at that time, taking the household's decision rule as given (Markov-perfect equilibrium), and does not commit to future spending. Second, physical capital is perfectly mobile, so K_t is pinned by no-arbitrage with the world rate r_t^* and is not a state variable; households trade a non-contingent international asset at that same rate. Third, adjusting G involves a convex resource cost $AC_t = (\Omega/2)(G_{t+1} - G_t)^2$, which discourages large year-to-year changes (Klein et al., 2008; Fernández-Villaverde et al., 2015; Bachmann and Bai, 2013b,a; Mendoza, 1991). The model does not include monetary frictions,

seigniorage, or price stickiness. Fiscal policy is always financed by real taxes. To close the open economy without a debt-elastic interest-rate premium (Schmitt-Grohé and Uribe, 2003), we use a constant domestic-debt rule and do not impose a tax ceiling on the baseline government. The baseline government directly maximizes household welfare; a smooth convex cost of tax deviations is added only as a robustness check. Government spending is also limited by an expenditure ceiling $G_{t+1} \leq \bar{g}\, Y_{t+1}$. This is the model counterpart of budget rules that cap next year's appropriations relative to projected output (see Section 4.3); here Y_{t+1} is known at t. These restrictions are consistent with the fiscal-rule designs discussed in Ghosh et al. (2013); Holm-Hadulla et al. (2012); Heinemann et al. (2018); Caselli and Reynaud (2020). The constant-debt closure also matches the debt-sustainability findings of Bohn (1998), who shows that primary surpluses systematically adjust to stabilize the debt stock.

The specification has no nominal rigidities and no household heterogeneity. With a single representative household and a real environment, the only intertemporal margins active in the model are those that matter for the welfare comparison we make: the path of disposable income (through τ_ℓ), the path of public consumption (through G), and the path of private saving (through A). Anticipated changes in fundamentals propagate through the user cost of capital, absorption, and external accounts in the standard way (Mendoza, 1991; Obstfeld and Rogoff, 1996; Schmitt-Grohé and Uribe, 2003), and one-period-ahead news on (z, r^*) is grafted onto the AR(1) structure following the news-shock literature (Beaudry and Portier, 2006; Jaimovich and Rebelo, 2009; Barsky and Sims, 2011; Schmitt-Grohé and Uribe, 2012; Beaudry et al., 2011; Kamber et al., 2017).

Table 2: Notation (selected symbols used in the model)

Symbol	Meaning
$\bar{B}$	Constant public debt level
K_t	Physical capital
A_t	Household asset holdings (claims on capital and bonds)
$\underline{a}$	Borrowing limit on private assets
TB_t	Trade balance (goods)
Y_t	Output; resource identity $Y_t = C_t + I_t + G_t + AC_t + TB_t$ (Eq. (17))
C_t	Private consumption
I_t	Investment
G_t	Government purchases (public spending flow)
AC_t	Adjustment costs entering absorption
r_t^*	World real interest rate
z_t	Total factor productivity (TFP)
θ_t	Preference weight on private consumption in $u(C_t, G_t)$

All variables are real and measured at the model's base frequency. In this document, "≡" means definitions and national-accounts identities, while "=" shows equilibrium conditions.

4.1 Households

Preferences. Household preferences are log–Cobb–Douglas in private and public goods:

$$u(C_t, G_t) \;=\; \theta_t \log C_t \;+\; \big(1-\theta_t\big)\,\log G_t, \qquad \theta_t \in (0,1). \tag{1}$$

A higher value of θ_t means households value private consumption more. When θ_t is lower, public goods become more valuable at the margin. The log specification ensures that both C_t and G_t are positive in equilibrium, so utility is always well defined. The process θ_t switches between two values, $\{\theta_{\text{low}}, \theta_{\text{high}}\}$, according to a symmetric two-state Markov chain:

$$\Theta \;=\; \begin{bmatrix} \rho_\theta & 1-\rho_\theta \\ 1-\rho_\theta & \rho_\theta \end{bmatrix}.$$

This process is orthogonal to TFP and the world interest rate (Bachmann and Bai, 2013b). We denote by $\bar{\theta} \equiv \mathbb{E}[\theta_t]$ the stationary mean.

Budget constraint and the Euler condition. The asset A_t covers both capital and bonds. Since no-arbitrage means that capital and bonds earn the same net returns, and the asset-income tax τ_a applies to both, A_t shows the household's financial position over time.

With a labor-income tax $\tau_{\ell,t}$ and an asset-income tax τ_a, the period budget constraint is

$$C_t \;+\; A_{t+1} \;=\; (1-\tau_{\ell,t})\, w_t L \;+\; \big[1+(1-\tau_a)\, r_t^*\big]\, A_t. \tag{2}$$

$$A_{t+1} \geq \underline{a}, \tag{3}$$

In this context, $\underline{a}$ represents a borrowing limit that is set outside the model. Because capital can move freely, K_t is determined by the firm's first-order condition (7), which depends on the current fundamentals (z_t, r_t^*). This means K_t is not an independent state variable. The household's asset position A_t is usually different from K_t and from $\bar{B}$, since A_t also includes net bond holdings.

Let λ_t be the multiplier on (2), where $\lambda_t = \theta_t / C_t$. If the household treats all assets in the same way, the Euler conditions become a single equation:

$$\frac{\theta_t}{C_t} \;=\; \beta\, \mathbb{E}_t\!\left[\frac{\theta_{t+1}}{C_{t+1}}\left(1+(1-\tau_a)\, r_{t+1}^*\right)\right], \tag{4}$$

This equation holds with equality when (3) is not binding. If the borrowing limit is binding, the equation becomes an inequality with complementary slackness.

If r_{t+1}^* is known one period ahead, then the after-tax world return on saving from t to $t+1$ is also known at time t. Still, when deciding on A_{t+1}, the household faces uncertainty, including the preference shock θ_{t+1} and new information arriving at date $t+1$. These factors affect the government's choice of G_{t+2}, which, through adjustment costs and the budget identity (15), influences the tax rate $\tau_{\ell,t+1}$ and next period's consumption.

4.2 Firms and shocks

Technology, factor prices, and capital mobility. Output is given by a Cobb-Douglas production function: $Y_t = z_t K_t^{\alpha} L^{1-\alpha}$, where α is between 0 and 1. In this case, the labor supply L stays constant.

$$Y_t \;=\; z_t \, K_t^{\alpha} L^{\,1-\alpha}. \tag{5}$$

In competitive factor markets, the following conditions hold:

$$r_{k,t} \;=\; \alpha \, z_t \left(\frac{K_t}{L} \right)^{\alpha-1}, \qquad w_t \;=\; (1-\alpha) \, z_t \left(\frac{K_t}{L} \right)^{\alpha}. \tag{6}$$

With perfect capital mobility, the rental rate after accounting for depreciation matches the world interest rate, $r_{k,t} - \delta = r_t^*$. This leads to the following static capital demand:

$$\alpha \, z_t \left(\frac{K_t}{L} \right)^{\alpha-1} \;=\; r_t^* + \delta. \tag{7}$$

If we solve equation (7) for K_t, we find that, in per-worker terms, $k_t \equiv K_t/L = (\alpha z_t/(r_t^* + \delta))^{1/(1-\alpha)}$. Alternatively, we can write:

$$K_t \;=\; L \left(\frac{\alpha \, z_t}{r_t^* + \delta} \right)^{\frac{1}{1-\alpha}}. \tag{8}$$

If information about (z_{t+1}, r_{t+1}^*) is available one period in advance, then at time t firms choose the following target capital stock for the next period:

$$K_{t+1} \;=\; L \left(\frac{\alpha \, z_{t+1}}{r_{t+1}^* + \delta} \right)^{\frac{1}{1-\alpha}}. \tag{9}$$

Current investment, $I_t \equiv K_{t+1} - (1-\delta) K_t$, reacts directly to new information about (z_{t+1}, r_{t+1}^*). In particular, K_{t+1} goes up when z_{t+1} increases and falls when r_{t+1}^* rises, as shown in standard anticipation experiments (Beaudry and Portier, 2006; Jaimovich and Rebelo, 2009).[5] If we plug (8) into the wage equation from (6), we get this reduced-form factor price:

$$w_t \;=\; (1-\alpha) \, z_t \left(\frac{\alpha \, z_t}{r_t^* + \delta} \right)^{\frac{\alpha}{1-\alpha}}, \tag{10}$$

This wage depends only on (z_t, r_t^*) and does not change with the asset position A_t.

Stochastic environment with news. The economy is driven by two outside factors. Changes in these factors are revealed at time t, but they affect the fundamentals at time $t+1$ (this is called

[5] With convex installation costs or time-to-build, (9) gives the target stock, not the actual stock within the period; the comparative statics are unchanged.

one-period-ahead news). Total factor productivity (TFP) follows a persistent AR(1) process:

$$\ln z_{t+1} = (1-\rho_z)\,\ln \bar{z} + \rho_z \ln z_t + \varepsilon_t^z, \qquad \varepsilon_t^z \sim \mathcal{N}(0, \sigma_z^2), \tag{11}$$

Here, $0 < \rho_z < 1$. We make the process discrete by using a three-state Rouwenhorst chain, which preserves its calibrated persistence. For the world interest rate, we use a two-state process to capture rare periods when financing is expensive.

$$r^*_{t+1} \;=\; \begin{cases} r^{*,H} = 5.52\%, & \text{with probability } p_H = 0.075, \\ r^{*,L} = 2.48\%, & \text{with probability } 1 - p_H, \end{cases} \qquad \text{i.i.d. across } t. \tag{12}$$

Most of the time, financing conditions are normal, but about once every thirteen periods, they shift to an expensive regime. With the news timing, agents get a one-year warning. The stationary mean is 2.71%. This setup is a deliberate choice: we use the rare-episode arrival structure and keep the AR(1)-implied node spread to calibrate severity. Specifically, the nodes are set at $\bar{r}^* \pm \sigma_r/\sqrt{1-\rho_r^2}$ with $(\bar{r}^*, \rho_r, \sigma_r) = (0.04, 0.85, 0.008)$. This means that persistence and volatility estimates for emerging-market world rates (Neumeyer and Perri, 2005; Uribe and Yue, 2006) affect only the severity of an episode. The arrival process is independent over time and does not inherit other AR(1) features. We make this change because the main fiscal question is how the government responds to sudden, rare tightenings of external financing. A two-state rare-episode chain captures this directly, while a smoothly mean-reverting rate process does not. This approach follows the sudden-stop events described by Calvo et al. (2004). The reference value $\bar{r}^* = 4\%$ is the midpoint between nodes. Although the chain never actually visits this state, it is used for reference accounting and Lemma 1. The lemma's world-rate elasticity is anchored there. At the modal node $r^{*,L} = 2.48\%$, the elasticity is about -10 instead of -7.1. We use ε_t^r for the draw at time t from (12), which is the news innovation revealed at t. At time t, the information set includes $\{z_t, r_t^*, \varepsilon_t^z, \varepsilon_t^r\}$, so the realized values are (z_{t+1}, r^*_{t+1}). In this paper, high-r^* states are referred to as rare expensive-financing episodes, rather than as phases of a persistent rate cycle. This is the standard news-shock timing (Beaudry and Portier, 2006; Jaimovich and Rebelo, 2009; Barsky and Sims, 2011; Schmitt-Grohé and Uribe, 2012), which makes one-period-ahead returns predictable when investment is chosen. We use the polar case where the full one-period-ahead innovation is anticipated, while the empirical literature often finds only partial anticipation. This timing matches the fiscal institution in the model. In the Philippines and most emerging markets, annual budgets are set a year in advance based on published macroeconomic projections for output and interest rates. In the model the counterpart is exact: (z_{t+1}, r^*_{t+1}) are known at t, so Y_{t+1} is known when the government chooses G_{t+1}. Two features of the equilibrium limit the quantitative impact of the polar assumption. The government's purchases rule depends on the current state $(A_t, G_t, z_t, r_t^*, \theta_t)$ but not on the news. This is a restriction from the equilibrium policy class, and its effect is measured in Section 5.1. Anticipation affects the household's saving decision, which uses the known (z_{t+1}, r^*_{t+1}), and the timing that makes Y_{t+1} known when the next budget is set. So, the full-anticipation timing is an

assumption about what the household knows and the budget calendar, not an extra choice in the government's strategy.

4.3 Government and fiscal financing

Budget constraint and the constant-debt rule. The government buys public goods, represented by G_t, and holds one-period domestic bonds, B_t, which pay an interest rate $r_{b,t}$. Because of no-arbitrage with the world asset market, $r_{b,t}$ equals r_t^*, as explained in the timing discussion below equation (22). The government's flow budget constraint is:

$$B_{t+1} \;=\; (1 + r_{b,t})\, B_t \;+\; G_t \;+\; AC_t \;-\; T_t \tag{13}$$

$$T_t \equiv \tau_{\ell,t}\, w_t L \;+\; \tau_a\, r_{b,t}\, A_t, \tag{14}$$

Here, revenue T_t comes from a proportional labor-income tax $\tau_{\ell,t} \geq 0$ and a proportional tax τ_a on private asset income. Policy keeps the bond stock constant:

$$B_{t+1} = B_t \equiv \bar{B}.$$

This setup is similar to the debt anchors used in many emerging-market fiscal frameworks during normal times (Schaechter et al., 2012; Bova et al., 2014). It separates the spending and financing decisions from debt issuance, which would otherwise help absorb economic cycles. Applying the rule in (13) means the government must run a primary surplus each period equal to the interest on the outstanding debt, $T_t = G_t + AC_t + r_{b,t}\bar{B}$. Solving for the labor tax that achieves this gives:

$$\tau_{\ell,t} \;=\; \frac{G_t \;+\; AC_t \;+\; r_{b,t}\,\bar{B} \;-\; \tau_a\, r_{b,t}\, A_t}{w_t\, L}, \qquad 0 \leq \tau_{\ell,t} < 1. \tag{15}$$

Changes in G affect households' disposable income as follows: when the debt stock remains constant, any change in G_t or AC_t is balanced by adjustments in $\tau_{\ell,t}$ (Baxter and King, 1993; Leeper, 1991). There is no strict upper limit on the tax rate below one, and the baseline government faces no separate penalty on tax deviations. Their cost is evaluated through household welfare. In the computational model, the tax rate stays within $[0, 0.99]$. Section 5 shows the long-run distribution of the tax rate and how often it exceeds 0.50.

Interpretation of τ_ℓ as a financing wedge. In this model, labor supply L does not change in response to taxes, so τ_ℓ does not create the usual labor-supply tax distortion because it does not affect how much labor is supplied. However, τ_ℓ still causes two important distortions for welfare. First, when τ_ℓ changes over time, it changes disposable income $(1 - \tau_{\ell,t})w_t L$ each period. This affects whether households choose to consume or save, as shown by the Euler equation (4). Second, because international asset markets are incomplete and the household can trade only a single non-

contingent asset, it cannot protect itself from extra fluctuations in disposable income caused by a changing τ_ℓ. This means τ_ℓ also distorts the risk-bearing margin. The borrowing limit (3) exists in the model but does not bind in the ergodic simulation (see Section 5.1), so the main issue is the lack of state-contingent assets, not the borrowing limit. The welfare impact of τ_ℓ comes from (i) its effect on the choice between consumption and saving through disposable income, (ii) its effect on risk-bearing when markets are incomplete, and (iii) the (C, G) allocation that affects utility through (1). In this paper, we refer to τ_ℓ as a residual labor-income financing wedge that balances the period budget under the constant-debt rule, rather than a tax on labor supply, but it still creates distortions for welfare in these ways. If labor were elastic, there would be an extra distortion from labor supply. The financing-assignment result in Section 6.3 would then measure the welfare cost from sources other than labor supply. We discuss this as a possible next step in the conclusion.

Adjustment costs. Changing G is costly:

$$AC_t \;=\; \frac{\Omega}{2}\left(G_{t+1} - G_t\right)^2, \qquad \Omega > 0. \tag{16}$$

The cost represents a straightforward use of resources. It shows up in the government budget and the goods-market identity, but not in utility, as described by Bachmann and Bai (2013b). This helps maintain the persistence of G that the data require, without creating intertemporal preference wedges.

The government's constraints are set by the constant-debt rule, the budget identity (15), the adjustment cost, and the spending ceiling $G_{t+1} \leq \bar{g}\, Y_{t+1}$. Here, Y_{t+1} is known at time t under the news timing (see Section 4.5). Section 5 shows the actual G/Y distribution and how often the ceiling is binding.

4.4 Market clearing and balance of payments

The goods market reaches equilibrium through the absorption identity:

$$Y_t \;=\; C_t + I_t + G_t + AC_t + TB_t, \tag{17}$$

Here, investment is defined as capital accumulation:

$$I_t \;=\; K_{t+1} - (1-\delta)\, K_t, \tag{18}$$

The trade balance is determined as the residual:

$$TB_t \;=\; Y_t - \big(C_t + I_t + G_t + AC_t\big). \tag{19}$$

When TB_t is positive, the country has net exports. If it is negative, the country is borrowing from abroad. The external position of the economy is included in the bond component of A_t.

Households can hold both domestic government bonds and foreign bonds at the world interest rate, but net foreign assets are not considered a separate state variable. This method follows the standard absorption model (Obstfeld and Rogoff, 1996; Corsetti and Müller, 2013). The capital accumulation equation, $K_{t+1} = (1-\delta)K_t + I_t$, is just a rearrangement of (18):

$$K_{t+1} \;=\; (1-\delta)\,K_t + I_t. \tag{20}$$

Timing and information. At the start of period t, agents see the state vector

$$\mathbf{s}_t \;\equiv\; \big(A_t,\, G_t,\, r_t^*,\, z_t,\, \varepsilon_{r,t},\, \varepsilon_{z,t},\, \theta_t\big).$$

Here, $\varepsilon_{r,t} \equiv \varepsilon_t^r$ and $\varepsilon_{z,t} \equiv \varepsilon_t^z$ are the news innovations in (12) and (11), which reveal (r_{t+1}^*, z_{t+1}) at time t. The sequence within each period is as follows: (i) firms set K_{t+1} using (9), based on the known values of (z_{t+1}, r_{t+1}^*); (ii) the government selects $G_{t+1} \in \Gamma(\mathbf{s}_t)$ and pays the adjustment cost AC_t from (16), with the labor tax $\tau_{\ell,t}$ determined by (15); (iii) the household, after seeing G_{t+1} and $\tau_{\ell,t}$, chooses $\{C_t, A_{t+1}\}$ subject to (2) and (3), and the Euler condition (4) holds given the after-tax world return implied by $r_{b,t} = r_t^*$, as set by the no-arbitrage condition (22); (iv) goods-market clearing (17) is achieved. The state then updates to

$$\mathbf{s}_{t+1} \;\equiv\; \big(A_{t+1},\, G_{t+1},\, r_{t+1}^*,\, z_{t+1},\, \varepsilon_{r,t+1},\, \varepsilon_{z,t+1},\, \theta_{t+1}\big).$$

At time $t+1$, the new variables (r_{t+2}^*, z_{t+2}) become known. Given the constant-debt rule and $r_{b,t} = r_t^*$, the period-t sequential budget identity is

$$T_t \;\equiv\; G_t + AC_t + r_t^*\,\bar{B} \;=\; \tau_{\ell,t}\,w_t L \;+\; \tau_a\,r_t^*\,A_t. \tag{21}$$

No-arbitrage between capital and bonds implies

$$r_{k,t} - \delta \;=\; r_{b,t} \;=\; r_t^*. \tag{22}$$

4.5 Markov-perfect equilibrium

In each period t, the government chooses next-period purchases $G' \equiv G_{t+1}$, and the household chooses next-period assets $A' \equiv A_{t+1}$. Because the government cannot commit to a plan for G that depends on future states, it takes as given that its future selves' decisions will depend only on the payoff-relevant state then prevailing. We seek a Markov-perfect equilibrium (Maskin and Tirole, 2001) in which both players make choices based only on the current state.

State, feasibility, and household best response. The Markov state is

$$\mathbf{s}_t \;\equiv\; (A_t,\, G_t,\, r_t^*,\, z_t,\, \varepsilon_t^r,\, \varepsilon_t^z,\, \theta_t).$$

Here, $(\varepsilon_t^r, \varepsilon_t^z)$ pin down (r_{t+1}^*, z_{t+1}) through (12) and (11). Given $(\mathbf{s}, G')$, the adjustment cost is $AC(\mathbf{s}, G') \equiv (\Omega/2)(G' - G)^2$, and the implementing labor tax follows from (15):

$$\tau_\ell(\mathbf{s}, G') \;\equiv\; \frac{G + AC(\mathbf{s}, G') + r^* \bar{B} - \tau_a r^* A}{w(z, r^*)\, L}.$$

The government's feasible set is

$$\Gamma(\mathbf{s}) \;\equiv\; \Big\{ G' : 0 < G' \leq \bar{g}\, Y_{t+1}(\mathbf{s}) \Big\}. \tag{23}$$

Purchases must be positive and cannot exceed $\bar{g}\, Y_{t+1}(\mathbf{s})$. Next-period output is determined by (9) and is known at time t because of the information timing. The constraint is the model counterpart of budget rules that cap next year's appropriations relative to projected output; here the projection is exact because Y_{t+1} is known. The ceiling matters in equilibrium: it restricts the government's discrete best response in about 38% of the states visited (see Section 5 for details and the realized G/Y distribution). When spending demand is high, purchases track Y_{t+1} up to the ceiling, so the procyclicality of purchases in equilibrium comes from both the government's optimization and this rule.

The lower-case a represents the asset position of an individual household. The upper-case A in the aggregate state $\mathbf{s}$ stands for the economy-wide representative agent's asset position, which is what policies are based on. In equilibrium, a equals A. When the government selects a possible policy $G' \in \Gamma(\mathbf{s})$ and there is a continuation value function V, the household decides on (c, a') to solve

$$J(a, \mathbf{s}, G'; \Psi, H) = \max_{c,\, a' \geq \underline{a}} \Big\{ \theta \log c + (1 - \theta) \log G + \beta\, \mathbb{E}\big[V(a', \mathbf{s}'; \Psi, H) \big] \Big\}. \tag{24}$$

This problem is subject to (2), with the household's individual asset a replacing aggregate A and with $\tau_{\ell,t} = \tau_\ell(\mathbf{s}, G')$, as well as the borrowing limit in (3). The next-period aggregate state is $\mathbf{s}' = (A', G', r_{t+1}^*, z_{t+1}, \varepsilon^{r\prime}, \varepsilon^{z\prime}, \theta')$, and the expectation averages over the next news innovations and preference draw at time $t+1$. Let $a' = h(a, \mathbf{s}, G')$ be the optimal asset choice. The representative-agent transition is $A' = H(\mathbf{s}, G') \equiv h(A, \mathbf{s}, G')$, and consumption is determined by the budget constraint.

Government problem and equilibrium definition. Taking the household best-response mapping H and the continuation value V as given, the time-t government chooses

$$\Psi(\mathbf{s}) \;=\; \arg \max_{G' \in \Gamma(\mathbf{s})} J\big(A, \mathbf{s}, G'; \Psi, H\big). \tag{25}$$

The government directly maximizes the household's value. Thus, the baseline is the discretionary household-welfare optimum subject to the fiscal constraints described above. A *Markov-perfect equilibrium* is a triple (Ψ, h, V) such that (i) $V(a, \mathbf{s}) = J\big(a, \mathbf{s}, \Psi(\mathbf{s}); \Psi, H\big)$ for all feasible $(a, \mathbf{s})$; (ii) $h(a, \mathbf{s}, G')$ solves the household problem (24) for each feasible $(a, \mathbf{s}, G')$, with $H(\mathbf{s}, G') \equiv h(A, \mathbf{s}, G')$

under the representative-agent restriction; (iii) Ψ solves the government's best-response problem (25); and (iv) the firm's optimality conditions, the constant-debt budget identity, and goods-market clearing all hold given the induced allocations and prices. We use $h(\cdot,\cdot,\cdot)$ for the household's off-path best response, defined for any feasible G', and $H(\cdot,\cdot)$ for the corresponding aggregate transition; on the equilibrium path, $G' = \Psi(\mathbf{s})$.

Tax-deviation robustness specification. As a robustness check, we also add a smooth convex term to the government objective,

$$P(\mathbf{s}, G') \;\equiv\; \frac{\lambda_\tau}{2}\,\mathbb{E}\Big[\big(\tau_\ell(\mathbf{s}', G') - \tau_{\ell,\text{ref}}\big)^2 \,\Big|\, \mathbf{s}\Big], \qquad \lambda_\tau = 0.50, \tag{26}$$

and replace the objective in (25) with $J - P$. This term is a reduced-form stand-in for collection, compliance, administrative costs, and labor-supply distortions that are not included under inelastic labor. It enters only the government objective: it is not the real adjustment cost AC_t, does not use resources, and does not enter household utility. Section 6.3 compares this robustness specification with the no-penalty baseline.

4.6 The financing-assignment identity

The constant-debt rule (15) links the cyclical patterns of τ_ℓ and G. Choosing one determines the other. The following lemma summarizes this connection in a single analytical result, evaluated at $A = 0$. Section 5 explains that this is the neighborhood of the ergodic asset distribution, not the paper's reference-point endowment, and that the identity is an impact relation rather than a fixed point of the household Euler equation. This lemma sets up the welfare analysis in Section 6.3. When comparing cyclical fiscal-rule reforms, the main policy question is how to split cyclical adjustment between G and τ_ℓ, instead of focusing on just one.

Lemma 1 (Duality of cyclical fiscal stance under residual financing)**.** *Consider a perturbation of the deterministic reference point.*

Here, $\varepsilon_G \equiv d\log G_t / d\log Y_t$, $\varepsilon_\tau \equiv d\log \tau_{\ell,t} / d\log Y_t$, $\varepsilon_w \equiv d\log w_t / d\log Y_t$, *and* $\varepsilon_r \equiv d\log r_t^* / d\log Y_t$ *show how government purchases, the labor tax, the wage, and the world rate change with output. All elasticities are measured at the deterministic reference point.*

Under the constant-debt rule and the period budget identity (15)*, evaluated at the deterministic reference point where* $A_{\text{ref}} = 0$*, the second equality for* ω *holds because* $\tau_a r^* A_{\text{ref}} = 0$ *at this point:*

$$\begin{aligned} \varepsilon_\tau \;&=\; \omega\,\varepsilon_G - \varepsilon_w + (1-\omega)\,\varepsilon_r, \\ \omega \;&\equiv\; \frac{G_{\text{ref}}}{G_{\text{ref}} + r^*\bar{B}} \;=\; \frac{G_{\text{ref}}}{\tau_{\ell,\text{ref}} w_{\text{ref}} L}. \end{aligned} \tag{27}$$

In a small open economy where capital is determined by the static no-arbitrage condition (8)*,* $\varepsilon_w = 1$ *for both productivity and world-rate shocks. For productivity shocks,* $\varepsilon_r = 0$*, so the duality becomes* $\varepsilon_\tau = \omega\,\varepsilon_G - 1$.

For world-rate shocks, $\varepsilon_r \neq 0$. The additional $(1-\omega)\varepsilon_r$ term reinforces the channel because $\varepsilon_r < 0$ for an r^* shock. This means that output moves in the opposite direction from r^* under the static no-arbitrage condition. In the no-debt-service limit $r^*\bar{B} \to 0$ (so $\omega \to 1$), the duality becomes

$$\varepsilon_\tau = \varepsilon_G - \varepsilon_w.$$

This result applies regardless of which shock causes the cyclical change.

Remark (Generalisation to debt-elastic financing rules). The duality identity does not rely on the constant-debt rule. Instead, suppose public debt follows a rule that changes with the economic cycle:

$$B_{t+1} \;=\; \bar{B} \cdot \big(1 + \varepsilon_B \cdot (\log Y_t - \log Y_{\text{ref}}) + \cdots\big), \qquad \varepsilon_B \in \mathbb{R}.$$

Here, ε_B measures how sensitive debt issuance is to the economic cycle. If $\varepsilon_B = 0$, this gives the constant-debt rule of Lemma 1. If $\varepsilon_B > 0$, debt increases during economic booms. If we repeat the reference-point derivation above and include this extra cyclical debt-issuance term in the period budget, we get

$$\varepsilon_\tau \;=\; \omega\,\varepsilon_G \;-\; \varepsilon_w \;+\; (1-\omega)\,\varepsilon_r \;-\; \frac{\bar{B}}{\tau_{\ell,\text{ref}}\, w_{\text{ref}}\, L}\,\varepsilon_B. \tag{28}$$

In the calibration, $\bar{B}/(\tau_{\ell,\text{ref}}\, w_{\text{ref}}\, L) \approx 2.01$, so each unit of cyclical debt-issuance elasticity lowers the implementing-tax elasticity ε_τ by about 2.01 units. The financing-assignment channel still operates for any ε_B short of the perfectly state-contingent value

$$\varepsilon_B^* \;=\; \frac{\omega\,\varepsilon_G - \varepsilon_w + (1-\omega)\,\varepsilon_r}{\bar{B}/(\tau_{\ell,\text{ref}}\, w_{\text{ref}}\, L)},$$

which sets ε_τ to zero. A debt rule with $\varepsilon_B = \varepsilon_B^*$ fully absorbs the cyclical financing residual for the shock in question. This is the single-shock analogue of the state-contingency that Lucas–Stokey complete markets provide: a single output-indexed issuance elasticity cannot replicate fully state-contingent debt when there are multiple shocks. In practice, both market discipline and rule design (Schaechter et al., 2012; Bova et al., 2014; Guerguil et al., 2017) keep the cyclical debt-issuance elasticity well below ε_B^* in emerging-market economies. Therefore, the financing-assignment channel matters under any debt-stabilizing fiscal rule, not just the constant-debt rule used in our experiment.

Scope of the lemma. There are two important scope qualifications. First, Lemma 1 is an *impact* identity. In the proof, $dA_t = 0$ because assets are predetermined within the period. This makes the result exact at the moment of impact, but not over a simulated path where A_t changes each period. The unconditional cyclical comovements shown in the simulation tables (for example, $\text{corr}(G, Y)$ and $\text{corr}(\tau_\ell, Y)$) reflect the sign of the identity, but not its precise size. Second, the elasticities ε_G, ε_τ, and ε_r are defined for each shock in a two-shock setting. Specifically, $\varepsilon_r \equiv d\log r^*/d\log Y$ is zero for a productivity shock and negative for a world-rate shock. Therefore, (27) should be interpreted

as a separate identity for each shock, as the lemma states, not as a single unconditional value.

This duality has three consequences for the welfare comparisons that follow. First, if a reform package sets an exogenous cyclical rule for G, it also sets a specific cyclical rule for τ_ℓ, based on (27). Thus, the welfare comparison is between two joint (G, τ_ℓ) packages, not just between two G-rules. Second, the slope of ε_τ with respect to ε_G is ω, which is fixed by the reference-point budget shares. The intercept depends on the shock-specific elasticities ε_w and ε_r. These are structural features of the constant-debt environment, not equilibrium outcomes. Third, in our calibration, $\omega = G_{\text{ref}}/(G_{\text{ref}} + r^*\bar{B}) = 0.2288/(0.2288 + 0.04 \times 0.5) \approx 0.92$. This means that, for productivity shocks, the implementing-tax elasticity changes a little less than one-for-one with the spending rule's elasticity. For world-rate shocks, $\varepsilon_r \approx -7.1$ in the calibration, and the extra term $(1-\omega)\,\varepsilon_r \approx -0.57$ makes ε_τ even more negative. This strengthens the financing-assignment channel when output is low and r^* is high. Section 6.3 shows the quantitative welfare effects of choosing one assignment over the other.

5 Calibration

All parameters are calibrated annually. Instead of using simulated method of moments, we calibrate by separating parameters set externally (using standard values from the literature or long-run definitions) from those chosen internally so that purchases are procyclical, the labor tax moves against output, and cross-border risk sharing remains limited. The HP cyclical moments in Table 5 are a check on that informal choice, not formal targets. Table 3 lists all parameters.

The externally set parameters follow standard practice. The discount factor $\beta = 0.96$ gives an annual real rate of about 4%. The capital share $\alpha = 0.33$ and depreciation rate $\delta = 0.10$ are typical values. The mean world rate $\bar{r}^* = 0.04$ matches the long-run benchmark used in SOE calibrations. The stationary mean preference weight $\bar{\theta} = 0.75$ places three times as much utility weight on private consumption as on public consumption, which is consistent with Baxter and King (1993); Bachmann and Bai (2013b). Stationarity of A_t in the open economy is achieved through strict impatience relative to the after-tax world return, so $\beta\left(1 + (1-\tau_a)\bar{r}^*\right) = 0.96 \times 1.034 \approx 0.993 < 1$, together with the borrowing limit (3) and a constant public debt level $\bar{B} = 0.5$. This is done without using a debt-elastic interest-rate premium. Because of strict impatience, the household Euler equation does not have an interior deterministic fixed point for A. The paper's reference point endows the household with claims on capital and public debt, $A_{\text{ref}} = K_{\text{ref}} + \bar{B}$, so the reference external position is balanced and the reference labor tax is 0.219 (Table 8). That point is a normalization of the asset-grid origin and a reference for scaling, not a steady state. The ergodic distribution of A_t sits near zero under either normalization, which is why Lemma 1 can be evaluated at $A = 0$. The alternative $A_{\text{ref}} = 0$ assigns those claims to foreigners and is reported as a robustness check in Appendix I. The constant-debt rule sets the interest-service requirement $r_t^*\bar{B}$ in the budget each period.

The expenditure ceiling is active in the ergodic simulation, and we report its distribution directly.

All figures use the unified tabular simulation pipeline on the 60×48 grid, with ten seeds and $T = 10{,}000$. Simulated G always stays within its numerical grid bounds. The expenditure ceiling $\bar{g}$ is enforced in the government's discrete best response (23). The best-response argmax is at the ceiling in about 38% of ergodically visited states. In areas with high spending demand, discrete best-response purchases follow Y_{t+1} along the constraint.

Along the simulated path, the realized ratio G_t/Y_t exceeds $\bar{g} = 0.35$ in about 27% of periods, with a maximum of 0.48. The converged policy uses damped, regularized polynomial updates of the discrete best response, and the fitted rule smooths over the ceiling kink instead of enforcing it in every state. These excursions occur because of approximation slack in the policy-function fit, not because of the environment. Adding the tax-deviation term in (26) as a robustness check changes the government's one-step discrete best response in only 1.4% of ergodically visited states and leaves the converged headline policy unchanged (see Section 6.3).

The tax margin is clearly strained in equilibrium. The ergodic implementing tax has a mean of 0.479 and a standard deviation of 0.072, goes above 0.50 in 27% of periods, and reaches a maximum of 0.74. Since there is no hard ceiling on the tax below one, these excursions are equilibrium outcomes, not constraint violations. The cost comes from the household's marginal-utility cost of financing, which also drives the fiscal-feasibility mechanism behind the state-conditional asymmetry in Section 6.3. An extra point of implementing tax reduces consumption in states where consumption is already low, so the effective marginal cost of tax-financed purchases is highest in low-productivity, high-world-rate states. Throughout the paper, "fiscal feasibility tightens" means this rising marginal-utility cost of financing, especially when the expenditure ceiling binds. The high ergodic mean comes from the model structure. Labor income, with a share of $1 - \alpha = 0.67$, is the effective tax base for financing an ergodic G/Y near 0.31 plus debt service, because the ergodic household asset position is close to zero whatever the reference point (Section 5), so asset-income revenue $\tau_a r^* A$ is second order.

The term $\tau_a r^* \mathbb{E}[A] \approx 0.0006$ is much smaller than labor-tax revenue. This means the ω approximation in Lemma 1, which sets $\tau_a r^* A = 0$, is accurate at the ergodic mean even though the reference point itself endows the household with claims on capital and public debt.

Two groups of parameters are set internally to match the emerging-market cyclical patterns described in Section 3. First, the adjustment cost $\Omega = 0.004$ controls how persistent G is through the resource cost (16). We choose this value so that G's persistence is in an empirically plausible range and $\sigma(G)$ does not become too large. The resulting autocorrelation in levels is 0.68. On the HP cyclical basis, the model's persistence is 0.19, which is below the data's 0.36 (see Section 6.2). Second, the preference-shock parameters $(\rho_\theta, \theta_{\text{low}}, \theta_{\text{high}}) = (0.90, 0.60, 0.90)$ account for the remaining variation in the spending share G/Y that productivity and the world rate cannot explain, following Bachmann and Bai (2013b). The symmetric two-state Markov chain on θ is set together with Ω so that the persistence and level volatility of G stay in an empirically plausible range. That is an informal internal choice, not a formal match to the HP cyclical moments in Table 5.

Table 3: Baseline calibration (annual frequency)

Parameter	Symbol	Value	Source / target
Preferences & technology			
Discount factor	β	0.96	Standard annual discount factor (implies $\approx 4\%$ real rate); see small-open economy calibrations such as Mendoza (1991).
Capital share	α	0.33	Standard capital share in growth/RBC models; consistent with cross-country evidence in Gollin (2002).
Depreciation	δ	0.10	Implies a 10% annual capital depreciation rate, in line with standard RBC practice (Mendoza, 1991).
Mean private-good weight	$\bar{\theta}$	0.75	Stationary mean of the preference weight θ_t in (1); chosen to allow a sizable role for public purchases, as in Baxter and King (1993); Bachmann and Bai (2013b).
International finance			
World-rate node midpoint	$\bar{r}^*$	0.04	Reference value anchoring the node placement and the reference-point accounting; the midpoint is never visited by the two-node chain, whose stationary mean is 2.71% (Section 4). Set to the conventional long-run world real rate in SOE models, $\approx 4\%$ annually (Mendoza, 1991; Schmitt-Grohé and Uribe, 2003).
Public debt (constant)	$\bar{B}$	0.50	Constant public debt *level* in output units (equal to 32.8% of reference-point output $Y_{\text{ref}} = 1.526$, at the low end of the 2000–2019 range for domestically held Philippine national-government debt), used to close the open economy without a debt-elastic interest-rate premium; cf. (Schmitt-Grohé and Uribe, 2003).
Asset-income tax	τ_a	0.15	Stylized effective tax rate on capital and bond income for an emerging economy; consistent with EM tax-structure discussions in Mueller et al. (2015); Vegh and Vuletin (2015); Talvi and Végh (2005).
Borrowing limit	$\underline{a}$	-1.50	Lower bound on private assets in (3); set equal to the lower bound of the asset grid.
Fiscal policy constraints and costs			
Adj. cost on ΔG	Ω	0.004	Small quadratic cost on changes in G to smooth fiscal dynamics, following the public-spending adjustment costs in Bachmann and Bai (2013b); matched internally to the AR(1) persistence of G at annual frequency.

Continued on next page

Parameter	Symbol	Value	Source / target
Expenditure ceiling	$\bar{g}$	0.35	Ceiling on G'/Y_{t+1} in the government's feasible set (23), the counterpart of appropriations-style caps on the year-ahead budget relative to next-period output; set above typical EM government-expenditure shares (World Bank, 2025a). Binds at the discrete best response at $\approx 38\%$ of ergodically visited states; the fitted policy's realized G_t/Y_t exceeds it in $\approx 27\%$ of periods (approximation slack; see text).
Tax-deviation robustness	λ_τ	0.50	Weight on the convex tax-deviation term (26) used only in the robustness specification; the baseline sets $\lambda_\tau = 0$. It is a reduced-form cost of moving the implementing tax away from $\tau_{\ell,\text{ref}}$ (collection, compliance, and omitted labor-supply distortions) and leaves the converged headline policy unchanged (Section 6.3).
Rule-target tax cap	$\bar{\tau}_\ell^{\max}$	0.50	Caps the *counterfactual rule targets* in Section 6.3 (Panel B). Not a constraint of the baseline government problem: the baseline ergodic implementing tax exceeds 0.50 in 27% of periods (see text). A stylized ceiling rather than an estimated Laffer peak.
Exogenous processes (annual)			
TFP AR(1)	(ρ_z, σ_z)	$(0.90, 0.03)$	Highly persistent and volatile TFP shocks to match emerging-market business-cycle facts (Aguiar and Gopinath, 2007); discretized with $n_z = 3$ Rouwenhorst states (Rouwenhorst, 1995; Kopecky and Suen, 2010).
World-rate process	(ρ_r, σ_r)	$(0.85, 0.008)$	Two-node rare expensive-financing episode process, eq. (12): the expensive regime occurs with probability 0.075 per period, independently over time (a sudden-stop-style arrival; cf. Calvo et al., 2004), with nodes $\bar{r}^* \pm \sigma_r/\sqrt{1-\rho_r^2} = \{2.48\%, 5.52\%\}$ whose spread is disciplined by the estimated AR(1) evidence (Neumeyer and Perri, 2005; Uribe and Yue, 2006); (ρ_r, σ_r) set episode severity (see Section 4).
Preference shock	ρ_θ	0.90	Persistent taste shocks shifting the weight on public vs. private consumption, as in Bachmann and Bai (2013b).
Preference states	$(\theta_{\text{low}}, \theta_{\text{high}})$	$(0.60, 0.90)$	Two-state preference process (low vs. high weight on private consumption), calibrated in the spirit of Baxter and King (1993); Bachmann and Bai (2013b).
News states	(z', r')	yes	One-period-ahead news on future TFP and world rate; agents observe (z_{t+1}, r^*_{t+1}) at time t, following news-shock frameworks in Beaudry and Portier (2006); Leeper et al. (2013).
Discretization & grids			

Continued on next page

Parameter	Symbol	Value	Source / target
Assets A	n_a	21–72	Asset grid on $[-1.5, 8]$ (in output units), with center-dense spacing near typical debt/asset levels, following numerical guidelines in (Judd et al., 2011). Headline grid $n_a = 60$.
Government G	n_g	21–60	Grid for government purchases with center-dense spacing around the reference-point value G_{ref} to accurately capture policy movements (Judd et al., 2011). Headline grid $n_g = 48$. Appendix I re-solves the model with the same number of nodes placed over the ergodic range of G instead.

Notes: All parameters are annual. There is no debt-elastic interest-rate premium (DEIR = 0). Stationarity is ensured by the borrowing limit and strict impatience relative to the after-tax world return ($\beta(1+(1-\tau_a)\bar{r}^*) < 1$), together with the constant public-debt rule $B_{t+1} = B_t \equiv \bar{B}$. The deterministic reference point implied by the calibration ($A_{\text{ref}} = K_{\text{ref}} + \bar{B}$; Table 8) is used to scale variables and to initialize the numerical solver and simulations described below. The headline grid for moments, welfare, and the policy-coefficient diagnostics is 60×48 (assets $\times$ spending); Appendix I reports the sensitivity of the results to how those nodes are placed. Agents receive one-period-ahead news on (z_{t+1}, r^*_{t+1}). Main modeling precedents for public spending and fiscal news/foresight include Bachmann and Bai (2013b), Beaudry and Portier (2006), and Leeper et al. (2013).

5.1 Computation

The Markov-perfect equilibrium is the fixed point of the household and government best-response operators on a discretized state $(A_t, G_t, z_t, r^*_t, \theta_t)$ together with one-period-ahead news (z_{t+1}, r^*_{t+1}). Appendix D records the iteration as Algorithm 1. Capital is not a discrete state: K_t is recovered from (z_t, r^*_t) by no-arbitrage. Moments and welfare use one 60×48 tabular pipeline. For the tax-deviation robustness specification the government block maximizes $J - P$ with $\lambda_\tau = 0.50$ and all other settings unchanged.

The purchases projection is grouped over the 12 current-state combinations (z, r^*, θ) and therefore does not depend on the news. That is a restriction of the policy class, not of the equilibrium. The household's tabular policies—the objects interpolated along the simulated path—do depend on the news. The restriction is small: with hysteresis off, the government's unrestricted one-step argmax changes across the six news nodes in a minority of cells, with an average range of 0.04 grid steps (maximum 9). Private agents look ahead; the purchases rule depends on current fundamentals.

5.2 Validation

After solving for the policy functions, we simulate $T = 10{,}000$ periods, discarding the first 500 as burn-in. All validation and moment tables use the unified tabular pipeline on the main 60×48 grid with ten random seeds, unless stated otherwise. Table 4 shows how the reference-point targets compare with the long-run simulation averages. By design, the deterministic reference point is $G_{\text{ref}}/Y_{\text{ref}} = 0.15$, but the ergodic mean of G/Y is higher, at 0.313 on the main grid. The simulated

path stays within the numerical grid bounds, and Section 5 describes the full G_t/Y_t distribution, including periods when it exceeds $\bar{g}$. This overshoot occurs because the Markov-perfect purchases rule is nonlinear in the state. The public-good preference state θ and changes in the world rate push the ergodic mean above the reference-point ratio. Both values are at the high end of the range for emerging markets and are above the Philippine purchases share of about 0.15. The broader expenditure-to-GDP ratio, which includes transfers, is about 0.20. We report this difference directly instead of adjusting the calibration. The welfare counterfactuals in Section 6.3 compare cyclical reassignments at a common long-run mean of G/Y, which isolates the effect of the assignment within the reported experiments; it does not establish invariance to recalibrating the level of G/Y.[6] We do not report simulated K/Y and I/Y because, with perfect capital mobility, K_t changes each period with (z_t, r_t^*) and I_t can be negative in some states. Their ergodic means are not calibration targets.

The core business-cycle comovements at annual frequency line up in sign with the data. Government spending co-moves positively with output—weakly on the cyclical basis ($\text{corr}(\tilde{G}, \tilde{Y}) \approx +0.08$) and more visibly in levels ($\text{corr}(G, Y) \approx +0.29$)—and the implementing labor tax is countercyclical ($\text{corr}(\tau_\ell, \tilde{Y}) \approx -0.33$). The signs are consistent with the Philippine data (Section 3) and with the cross-country emerging-market evidence. The reference-point ratios are the formal targets; Ω and the preference process keep the persistence and level volatility of G in a plausible range, but the HP cyclical moments in Table 5 are a check, not a fit. Section 6.2 quantifies where the magnitudes fall short of the data.

What the welfare comparison requires of the calibration. The welfare assessment compares scenarios using the same calibration, closure, shock paths, and baseline external position. Because the CEV is a proportional compensation for a log-utility household, scaling private consumption in every scenario by the same factor does not change the CEV. Holding the calibration fixed therefore isolates differences in the cyclical assignment across the reported scenarios. This does not mean that calibration levels are irrelevant: changing those levels can alter marginal-utility weights, policy functions, and the ergodic distribution.

Where the calibration matches the data and where it does not. The reference point matches its targets. Purchases are 15% of output, the external position is balanced, and the implementing labor tax is 21.9%, all consistent with Philippine magnitudes. The ergodic distribution the model then generates does not inherit those levels, and we report the discrepancies rather than adjust them away. The ergodic consumption share is $\mathbb{E}[C]/\mathbb{E}[Y] = 0.36$ against roughly 0.73 in

[6] A calibration targeting a lower ergodic mean of G/Y would require re-solving the model. It need not leave either the cyclical decomposition or welfare unchanged, because changing G_{ref} changes $\omega = G_{\text{ref}}/(G_{\text{ref}} + r^* \bar{B})$ in Lemma 1 and changes the marginal utility of public consumption. We retain $G_{\text{ref}}/Y_{\text{ref}} = 0.15$ as the reference-point target because it corresponds to government *purchases*—final consumption plus public investment, the national-accounts counterpart of the model's G—in the Philippine data. The broader ratio near 0.20 includes transfers and therefore is not the model-consistent target for G. The ergodic-mean overshoot is a property of the solved nonlinear equilibrium relative to this reference-point target.

Table 4: Targets vs. model (long-run averages)

Target	Reference point	Simulation mean
G/Y	0.150	0.313
K/Y	2.360	n.a.
I/Y	0.236	n.a.
$E[A]$	4.096	0.092

Notes: Reference-point values are from Table 8. Simulation means are computed under the unified tabular pipeline on the headline 60×48 grid ($T = 10{,}000$ periods, 500-period burn-in, ten seeds). The time-average G/Y exceeds the reference-point target due to the nonlinear procyclical purchases policy. Simulation means for K/Y and I/Y are omitted. In the small-open-economy specification, capital K_t is determined period by period by the no-arbitrage condition $F_K(K_t, 1) = r_t^* + \delta$ and thus varies with (z_t, r_t^*); accordingly, the ergodic mean of K/Y is not a calibration target. Investment is defined in the model as $I_t \equiv K_{t+1} - (1-\delta)K_t$ (cf. (20)) and is implied by the induced path of K_t; it can be negative in some states, so a simulation mean for I/Y is not targeted. The ergodic mean of A is strictly positive—the simulation never visits $A_t \leq 0$—with interquartile range $[0.01, 0.15]$ and 99th percentile 0.34, so the small-asset approximation $\tau_a r^* \mathbb{E}[A] \approx 0$ used for ω in Lemma 1 remains accurate at the ergodic mean ($\tau_a r^* \mathbb{E}[A] \approx 0.0006$, second order relative to labor-tax revenue). The gap between the reference value of A and its ergodic mean is the impatience effect discussed in Section 5: with $\beta\big(1 + (1-\tau_a)r^*\big) < 1$ the household runs its holdings down to near the borrowing region wherever the reference point is placed, so the ergodic mean is 0.092 here and 0.101 under the $A_{\text{ref}} = 0$ normalization even though the reference values differ by 4.096.

the data; the mean implementing tax is 0.479 and exceeds 0.50 in about a quarter of periods, well above measured effective labor-tax rates; and the ergodic external position is about -286% of output against an observed range of -10% to -20%.

It is worth being precise about why the last of these survives a reference point at which it is zero, because the two facts look contradictory. The reference asset level A_{ref} fixes the point at which the calibration is constructed; it does not pin the ergodic distribution of household assets, which is determined by the Euler equation on the asset grid. With $\beta\big(1 + (1-\tau_a)r^*\big) < 1$ the household is impatient relative to the after-tax world rate and decumulates toward the borrowing limit regardless of where the reference point is placed, so the ergodic mean of A is close to zero under either normalization. Since $\text{NFA}_t = A_t - K_t - \bar{B}$ and the country must finance a capital stock of 3.60 and public debt of 0.50, an ergodic mean household asset position near zero mechanically implies external liabilities of roughly the whole capital stock. Closing that gap requires a mechanism that makes the household want to hold domestic wealth—a higher discount factor, or a debt-elastic interest rate with a target the ergodic distribution is actually pulled toward—and not a different choice of A_{ref}. We flag this as the calibration's principal unresolved limitation. It bears on the magnitudes of every consumption-equivalent number in the paper, since those are read off marginal utility at the ergodic consumption level, and Appendix I shows that the comparisons themselves are unaffected by moving the reference point across the whole range this parameter spans.

The higher ergodic purchases share (0.31 compared with the 0.15 reference-point target) is a property of the solved equilibrium rather than of the normalization: it is the level distortion analyzed in Section 6.4, and moving the reference point leaves it essentially unchanged. Footnote 6 explains that the cyclical counterfactuals keep the long-run mean fixed and therefore do not, on their own, establish invariance to a level recalibration; Section 6.4 supplies the level experiment directly.

6 Results

Section 6.1 presents the fitted policy rules and the equilibrium-path crowding-out channel. Section 6.2 reports business-cycle moments, the model fit to emerging-market data, and the variance decomposition of G into a systematic component traceable to the macro state and a residual. Section 6.3 reports the welfare counterfactual: the 2×2 comparison of spending-side and financing-side assignments of the cyclical adjustment, with both panels on the same household-resolve footing.

6.1 Equilibrium policy rules

All numbers in this subsection are computed on the headline 60×48 grid, from the same converged solve that produces the moment and welfare tables in Sections 6.2–6.3.

Crowding out along the equilibrium path. The fitted household rule implies a negative elasticity of next-period assets with respect to current government purchases in every exogenous state group, with a cross-group mean of

$$\left.\frac{\partial \log A_{t+1}}{\partial \log G_t}\right|_{\bar{\mathbf{s}}} = -2.05,$$

so a 1% increase in G_t is associated with about a 2.1% decrease in A_{t+1} in the projection, holding the exogenous state fixed.[7] That slope is a local summary of a step-like tabular rule ($R^2 \approx 0.34$ in the projection fit) and is not quoted as a structural crowding-out elasticity: the simulation and welfare comparisons interpolate the tabular policy, not the projection. Because the coefficients are conditioned on (z, r^*, θ), and because the induced movement in G_{t+1} is negligible ($|\beta_1| \leq 0.016$ in Table 10), this is a within-state object rather than a correlation running through the fundamentals.

Panel A of Table 10 documents the pronounced heterogeneity of the asset-accumulation rule: persistence ranges from 0.86 to 2.03 and the response to current purchases from -0.73 to -3.49. Persistence exceeds one at the centering point in eleven of the twelve groups, which does not imply a divergent asset process: the fitted rule is concave in $\log A_t$ in every group, and the simulated path is generated by the tabular policy, whose ergodic distribution is stationary with $\mathbb{E}[A] = 0.09$ (Section 5). And because the projection fits a step-like rule with a smooth quadratic, the group-level coefficients should be read as slopes of a local summary rather than as structural estimates; the table note reports the fit statistics. Panel B shows that the government's purchases rule varies almost exclusively through the state-dependent level shift β_0; own persistence and asset sensitivity are small on average, and exactly zero in the high-productivity groups where the rule is constant in (A_t, G_t). Figure 1 shows both rules over the ergodic region: the saving rule shifts up with productivity and slopes down in current purchases (Panels A and B), while purchases are a flat,

[7] The fitted rule is $\log A_{t+1} = \alpha_0 + \alpha_1 \log A_t + \alpha_2 (\log A_t)^2 + \alpha_3 \log G_t + \alpha_4 (\log G_t)^2$; the elasticity at $\log G = g_{\text{norm}}$ equals the mean of α_3 across the 12 state groups. The regressors are centered at the simulated normalization constants a_{norm} and g_{norm}, which are ergodic-sample centering constants of the projection and are distinct from the deterministic reference point of Section 5.

state-contingent level (Panels C and D). The equilibrium purchases policy stays well inside the numerical bounds—its five values lie in $[0.229, 0.598]$ against a G-grid ceiling of 1.156, and no state selects a grid endpoint—so the heterogeneity in the rule is not driven by corner solutions. In short, household saving is state-contingent and crowded out by current G, while purchases are essentially a state-contingent intercept tracking (z, r^*, θ). The qualitative pattern is robust across (z, θ) and (r^*, θ) disaggregations: asset persistence is highest in low-TFP and high-r^* states, and the response to G is more negative in those same states. The full state-by-state coefficient tables are in the online supplement.

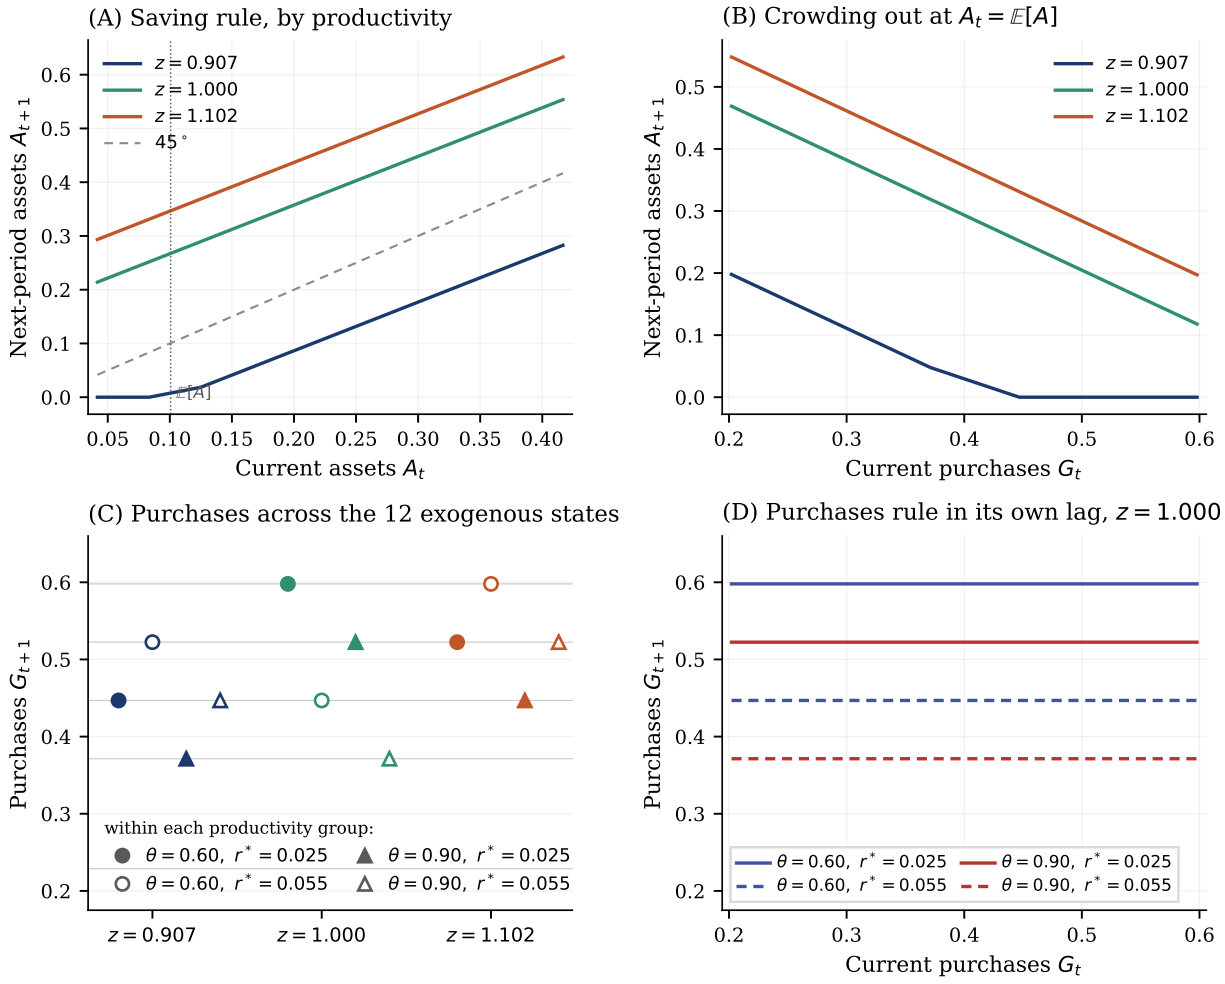


Figure 1: Equilibrium policy functions (A_{t+1}, G_{t+1}) across states.

Notes: Tabular equilibrium policies from the headline 60×48 solve—the objects the simulation interpolates, not the log-quadratic projections of Table 10—with the household block re-solved to a tolerance of 10^{-6} under the equilibrium g-policy. The 12 exogenous states are all combinations of productivity $z \in \{0.907, 1.000, 1.102\}$, world interest rate $r^* \in \{0.025, 0.055\}$, and preference weight $\theta \in \{0.60, 0.90\}$; all panels are drawn at the base news node. Panels A and C hold G_t at the median purchases node and panels B, C, and D hold A_t at its ergodic mean. Panels A and C are plotted over the ergodic support of assets, $[0, 1.25 \times A^{p99}]$, rather than the full computational grid $[-1.5, 8.0]$, in which the ergodic distribution occupies a narrow interval; simulated assets are never negative (Section 5). Panel A shows saving rising in productivity, with the low-productivity state below the 45° line. Panel B shows the crowding-out channel of Table 10, Panel A. Panel C plots all 12 states: the horizontal axis groups the three productivity levels, with colors as in Panels A and B, and within each group the marker shape denotes the preference weight and open markers the high world interest rate. Panels C and D share a common vertical scale: the equilibrium purchases rule is a state-contingent level taking only five distinct values over the entire state space, {0.229, 0.371, 0.447, 0.522, 0.598} (grey reference lines in Panel C), and is flat in its own lag.

6.2 Business cycle moments and the macro share of $\mathrm{Var}(G)$

A straightforward log–log regression of G_t on Y_t using the simulated data results in $R^2 \approx 0.12$ on levels ($R^2 \approx 0.007$ on HP cyclical components). Aggregate output therefore accounts for little of the cyclical variance of purchases. Preference (taste) dummies for θ account for about 2.7% of HP-cycled $\mathrm{Var}(\log G)$ on their own and about 2.6 percentage points incremental over productivity and world-rate dummies; jointly, (z, r^*, θ) account for about 7.5% of that cyclical variance (same baseline simulation). The residual is absorbed by the nonlinear policy, news, and sampling variation along the path.

Table 5 shows how the cyclical moments of G in the simulated economy compare to those in the Philippine data. Each model moment is calculated the same way as in the data, using HP $\lambda = 100$ cyclical components of logs. The model comes close to the relative volatility of purchases, with $\sigma(\tilde{G})/\sigma(\tilde{Y}) = 1.94$ in the model against 2.20 in the data. This is the main result emphasized in the emerging-market fiscal-cyclicality literature. It overstates both volatilities in absolute terms ($\sigma(\tilde{Y}) \approx 4.9\%$ versus 3.0%; $\sigma(\tilde{G}) \approx 9.6\%$ versus 6.6%), but by about the same factor, so their ratio stays close to the data's. The cyclical comovement is positive but weaker than in the data: $\mathrm{corr}(\tilde{G}, \tilde{Y}) = +0.08$ compared to $+0.34$. The model's procyclicality is more noticeable at lower frequencies: on raw levels, $\mathrm{corr}(G, Y) = +0.29$ (Table 11), since purchases follow the persistent parts of productivity and the world rate more than their year-to-year cyclical component. Cyclical persistence is lower in the model than in the data (0.19 versus 0.36 on the common HP basis). The levels-basis autocorrelation of 0.68 in Table 11 cannot be compared to the data value, so we do not use it to judge the fit. The model also shows another feature common in emerging markets, though not by design: consumption is more volatile than output ($\sigma(C)/\sigma(Y) > 1$ in coefficient-of-variation terms, Table 11), which matches the excess-consumption-volatility fact for emerging markets reported by Aguiar and Gopinath (2007). The match is directional rather than quantitative: the model's ratio exceeds typical emerging-market estimates, in line with its tendency to overstate absolute volatilities.

Relative volatility is a check rather than a formal target: the calibration pins the reference-point ratios in Table 4, and Ω and the preference process keep the persistence and level volatility of G in a plausible range. The HP moments in this table are then compared to the data, not fitted to them. The model comes closest on relative volatility. It understates year-on-year comovement and persistence: purchases are procyclical in sign and at lower frequencies, but the HP correlation is weaker than in the Philippine data. The welfare comparison that follows isolates the financing-assignment channel in that economy.

Measurement basis and cross-grid stability of the baseline moments. Two measurement conventions discipline the comparison. First, the model's comovement is basis-dependent in an economically informative way: on levels the comovement between purchases and output is +0.29 (+0.34 in logs)—close to the data's headline figure—while on the data-consistent HP cyclical basis it is +0.08. The model–data comparison in Table 5 is made entirely on the HP basis; levels-basis

Table 5: Business-cycle moments of cyclical G: model vs. data

Moment (cyclical G)	**Model**	**Data**
Std. dev. (%)	9.6	6.6
Relative volatility $\sigma(\tilde{G})/\sigma(\tilde{Y})$	1.94	2.20
AR(1)	0.19	0.36
$\text{corr}(\tilde{Y}, \tilde{G})$	0.08	0.34

Notes: All rows, model and data, are computed on HP($\lambda = 100$) cyclical components of logs; annual frequency; data are the Philippines 2000–2019 (Table 1). Relative volatility uses the model $\sigma(\tilde{Y}) = 4.9\%$ and the data value 3.0%. AR(1) = first-order autocorrelation of the HP cyclical component in both columns. Model moments are cross-seed means from the unified tabular pipeline on the headline 60×48 grid (ten seeds, $T = 10{,}000$, burn-in 500)—the same pipeline and grid as the welfare comparison in Table 6 (that table uses twenty seeds). Baseline-row cyclical moments in Table 6 use the same HP convention; the twenty-seed correlation there is +0.09 rather than the ten-seed +0.08 in this table. Levels-basis model moments (not comparable to the data rows) are reported in Table 11.

model moments appear separately in Table 11 so that the two bases are not conflated. Second, every moment is computed under the single unified simulation pipeline used for the welfare analysis. The cyclical moments are robust to how the state space is discretized, while the ergodic *level* is more sensitive: re-allocating the same number of purchases nodes over the ergodic range leaves $\text{corr}(\tilde{G}, \tilde{Y})$ at +0.08 but moves mean G/Y from 0.31 to 0.34, a drift whose consequence for the welfare comparison is discussed in Appendix I. The qualitative character of the baseline equilibrium—weakly procyclical, volatile purchases with an implementing tax whose correlation with output is negative—is unaffected.

6.3 Welfare counterfactual: the financing assignment

We compare the procyclical Markov-perfect benchmark with two types of fiscal-rule reforms. In the spending-side group, a cyclical rule is imposed on G and τ_ℓ adjusts to balance the budget. In the financing-side group, the rule is imposed on τ_ℓ and G adjusts instead. Both groups use the same constant-debt rule, the same tax-target cap where applicable, and a common target for the long-run mean of G/Y. The intended difference is how the cyclical adjustment is divided between G_t and $\tau_{\ell,t}$, as explained in Lemma 1; the one-pass scaling and grid discretization leave small realized mean-G/Y differences that we quantify below. We consider four rules in total, two from each group:

(A.i) Countercyclical G ($\eta_G = 0.5$). $G^{\text{cf}}_{t+1} = \kappa\, G_{\text{ref}}\, (Y_{\text{ref}}/Y_t)^{\eta_G}$ with $\eta_G = 0.5$ and κ chosen by a one-pass scaling so that $\mathbb{E}[G^{\text{cf}}/Y]$ equals the baseline mean. The labor tax $\tau^{\text{cf}}_{\ell,t}$ adjusts

residually through (15).

(A.ii) Constant G. $G_t = \bar{G}$ with $\bar{G}$ chosen so that $\mathbb{E}[G^{\text{cf}}/Y]$ equals the baseline mean. Removes all state-contingent variation in G.

(B.i) Output-indexed tax rule ($\eta_\tau = 0.5$). $\tau^{\text{cf}}_{\ell,t} = \tau^B_\ell \cdot (Y_t/Y_{\text{ref}})^{\eta_\tau}$ with $\eta_\tau = 0.5$ and the scale constant τ^B_ℓ chosen so that $\mathbb{E}[G^{\text{cf}}/Y]$ equals the baseline mean. The target is capped at $\bar{\tau}^{\max}_\ell = 0.50$ and floored at zero. Purchases adjust residually with the same one-period budget timing as the baseline: next-period purchases are set from current revenue at the current target,

$$G^{\text{cf}}_{t+1} = \tau^{\text{cf}}_{\ell,t}\, w_t L + \tau_a r^*_t A_t - r^*_t \bar{B},$$

and clipped at the numerical grid bounds (effectively at zero below). The expenditure ceiling $\bar{g}$ constrains the government's discretionary choice and is not re-imposed on the exogenous-rule path, whose residual purchases are pinned by the budget identity. This convention is symmetric with the baseline fitted policy, which also exceeds the ceiling in a share of periods because of approximation slack.

(B.ii) Constant τ_ℓ. $\tau_{\ell,t} = \tau^B_\ell$ with τ^B_ℓ chosen so that $\mathbb{E}[G^{\text{cf}}/Y]$ equals the baseline mean. Removes all state-contingent variation in the tax target; purchases adjust residually with the same timing as B.i.

In this section, each spending-side rule is named for the cyclical stance of the instrument it fixes. On the financing side we avoid the label "countercyclical τ_ℓ" for B.i because the tax *target* satisfies $\text{corr}(\text{target}, Y) = +1$ by construction: it is an *output-indexed tax rule*. A higher tax when output is high is the financing counterpart of spending restraint in a boom (see also Vegh and Vuletin, 2015). The effective implementing tax reported in the moment columns differs from that target and is typically negatively correlated with output, for the reasons explained below.

The two committed-rule cases (committed-G and committed-τ_ℓ, both fixed at the reference level) are not part of the comparison. Their CEVs mostly reflect changes in mean(G/Y) rather than the cyclical-assignment channel, so including them would distract from the main comparison.

Before examining the results, recall a key design feature from the introduction: output depends only on (z_t, r^*_t) through the static capital-demand condition, and labor supply is fixed. Thus, for a given shock sequence, Y_t follows the same path under all four rules. Countercyclical fiscal policy therefore cannot stabilize output in this environment. There is no demand channel, nominal block, or default-risk margin through which G can affect output. Any welfare cost of a countercyclical rule here comes from intertemporal allocation and risk sharing under incomplete markets, together with changes in the path of public consumption, rather than from labor-supply distortions, since labor is inelastic. The model is designed this way to isolate the financing-assignment channel. Holding other channels fixed, a negative CEV provides an upper bound on the net cost of the same rule in a richer environment where stabilization benefits could offset part of that cost.

Welfare is measured by the ergodic consumption-equivalent variation (CEV) λ that makes

households indifferent between the stationary benchmark and counterfactual allocations:

$$\log(1+\lambda) \;=\; \frac{\overline{u^{\mathrm{cf}}} - \overline{u^{\mathrm{mp}}}}{\overline{\theta}}, \tag{29}$$

where bars denote simulated ergodic means under the relevant allocation and $\overline{\theta}$ is the mean preference weight along the shared exogenous path. Equation (29) follows exactly from log utility: multiplying benchmark consumption by $(1+\lambda)$ adds $\theta_t \log(1+\lambda)$ to flow utility in each period. The compensation adjusts only private consumption, so the CEV is measured in private-consumption units, while the utility difference includes changes in both private and public consumption. A negative λ means that households prefer the Markov-perfect benchmark.

Welfare protocol. For each rule, we solve the household's EGM saving rule using the stochastic process for $(G_t, \tau_{\ell,t})$ that the rule creates and then run new simulations. Table 6 shows the mean CEV from 20 independent simulation seeds on the headline 60×48 grid, calculated with (29).[8]

The welfare comparison uses the ergodic-mean measure. This approach compares the benchmark's stationary allocation with the counterfactual's stationary allocation, so it excludes the transition from the benchmark to the counterfactual distribution. It is a long-run welfare comparison, not the welfare effect of switching to the new rule from the benchmark. Excluding the transition has no clear directional effect unless we know whether transition utility is above or below its counterfactual stationary value. If adopting the rule is costly, the ergodic measure tends to favor the reform. The advantage is that this method uses the entire simulated horizon T, rather than the shorter effective horizon $1/(1-\beta) \approx 25$ of a discounted comparison, which helps reduce simulation noise.

In (29), the denominator is the simulated mean $\overline{\theta}$ for each paired path, not the calibrated mean used as an approximation. This means that stochastic variation in θ_t and its relationship with consumption and purchases are retained in $\overline{u^{\mathrm{cf}}} - \overline{u^{\mathrm{mp}}}$. As a result, the CEV formula is exact for the simulated log-utility allocations.

The 2×2 comparison. Table 6 reports the household-resolved comparison on the headline 60×48 grid. Panel A is the spending-side family (rules A.i–A.ii), and Panel B is the financing-side family (rules B.i–B.ii). Both panels use the same household-resolve protocol and target the same long-run mean of G/Y; realized means differ modestly because the scaling is performed once and constant G must lie on the numerical grid. Subject to that qualification, the comparison isolates the assignment of cyclical adjustment between G_t and $\tau_{\ell,t}$. One terminological point deserves emphasis: the 2×2 is a *comparison of alternative policy packages*, not an additive decomposition. Rules A.i and B.i are two different joint (G, τ_ℓ) allocations; their CEVs do not sum to a common total, and neither is a "component" of the other.

[8] The household re-solve is a household-only fixed point, not a re-solved Markov-perfect equilibrium, because the government rule is exogenous.

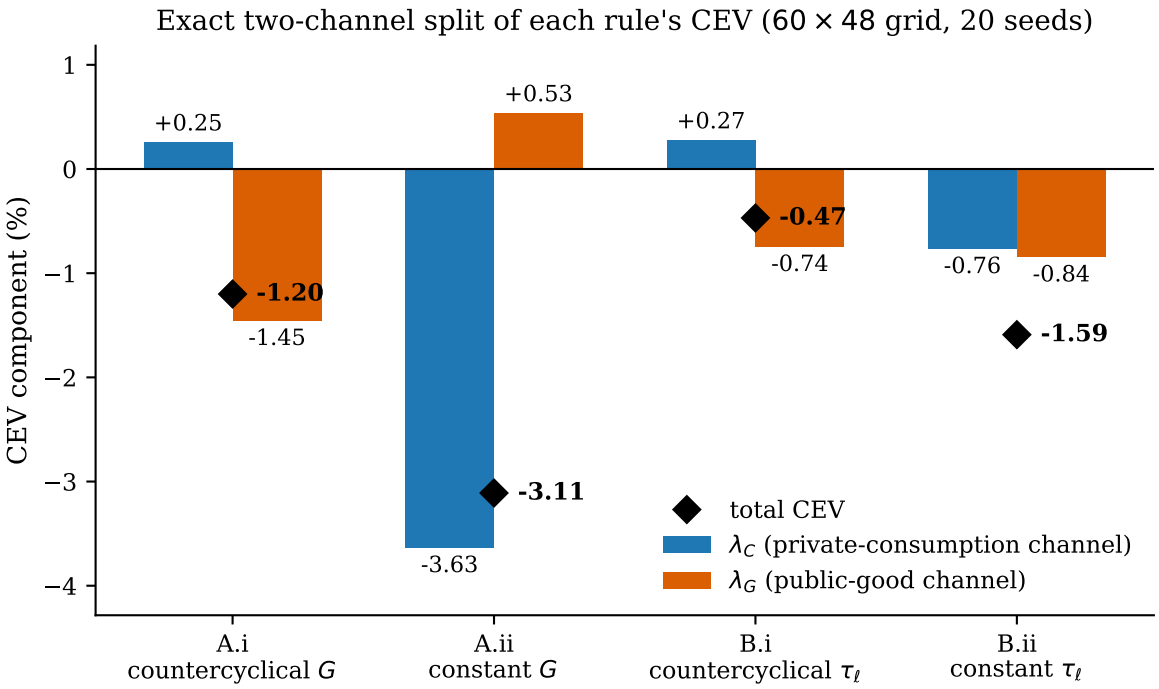


Figure 2: The mechanism exhibit: the exact multiplicative split $(1+\lambda) = (1+\lambda_C)(1+\lambda_G)$ of each rule's CEV into a private-consumption channel (λ_C) and a public-consumption channel (λ_G). On the headline grid every rule nets negative. The countercyclical purchases rule (A.i) destroys public-consumption value without delivering the financing relief (λ_C) that even the constant-purchases rule (A.ii) provides.

Notes: Values are the cross-seed means from Table 6 (headline 60 × 48 grid, household-resolve protocol, 20 seeds). Diamonds mark each rule's total CEV; bars show the two channels. Output volatility is invariant across all rules in this SOE because Y_t is determined by (z_t, r_t^*) through the static capital-demand condition rather than by the path of G_t.

What *can* be decomposed exactly is each rule's own CEV. Since utility is separable in $\log C$ and $\log G$, and each counterfactual uses the same exogenous paths as the benchmark, every CEV can be factored as $(1+\lambda) = (1+\lambda_C)(1+\lambda_G)$. Here, λ_C measures the effect of the rule on private consumption, which reflects the financing burden on the household, and λ_G measures the effect on public consumption. The last two columns of Table 6 report this split. Since labor supply does not change in response to taxes, a lump-sum levy timed the same way as the residual labor tax would affect household budgets in the same way. This means the λ_C column already shows how much of each rule's welfare effect comes from the way it is financed.

Reading the comparison. On the headline 60×48 grid every one of the four rules lowers welfare relative to the time-consistent benchmark. The spending-side rule (countercyclical G, with residual τ_ℓ) costs -1.20% in consumption-equivalent terms, its financing-side counterpart -0.47%, constant τ_ℓ -1.59%, and constant G -3.11%. Neither side of the budget is the cheap side: what separates the two assignments is not their average cost but the states in which they impose it.

Two considerations argue against reading the ordering of the four as a result. The first is the level protocol. The one-pass scaling that equalizes mean G/Y at the anchor seed leaves residual level mismatches that differ across rules, and iterating the anchors to a fixed point—so that each rule realizes the benchmark's spending share rather than merely targeting it—collapses the four cells into the band -1.08% to -1.86%, with constant G moving from -3.11% to -1.86%. Most of what separates the rules in the table is therefore level, not cycle. The second is the discretization. As Appendix C documents, the government's objective is flat enough in purchases that the median separation between the best node and the runner-up is worth 0.10% of consumption on the headline grid (0.02% after the node re-allocation), and re-allocating the same forty-eight nodes over the ergodic range changes the sign of three of the four cells. We therefore do not treat the ordering of the four ergodic cells as a result. On the headline grid they all lie below the benchmark; that sign is not robust to the node re-allocation. The computations support a band, not a ranking. What is separated, by an order of magnitude and with tight standard errors, is where each rule imposes its cost, which is the subject of the state-conditional decomposition below.

The split columns show where each cell comes from. The spending-side rule's loss is almost entirely a public-consumption loss and delivers essentially no financing relief ($\lambda_C = +0.25$, $\lambda_G = -1.45$). Forcing G to move against the cycle that the optimizing government prefers reduces public-consumption value while leaving the household's private-consumption path nearly unchanged from the benchmark. The countercyclical tax rule has the same shape with a smaller public-consumption concession ($\lambda_C = +0.27$, $\lambda_G = -0.74$): residual purchases retain some of the state-contingency that the Panel A rule legislates away, and public consumption carries an average utility weight of only 0.25 $(1-\bar{\theta})$. Constant τ_ℓ loses on both channels (-0.76 and -0.84). Constant G is the exception in the table, with a positive public-consumption term ($\lambda_G = +0.53$) and a large private-consumption loss ($\lambda_C = -3.63$), and it is the one cell whose sign is an artifact of its level rather than its cyclicality.

Table 6: The financing assignment: spending-side rules versus financing-side rules

	$\sigma(\tilde{Y})$ (%)	$\sigma(G)$	$\sigma(\tau_\ell)$	$\text{corr}(G,Y)$	$\text{corr}(\tau_\ell,Y)$	CEV (%)	λ_C (%)	λ_G (%)
Baseline (MPE)	4.94	0.074	0.072	+0.09	−0.33	—	—	—
Panel A: Smooth G, τ_ℓ residual								
Countercyclical G ($\eta_G = 0.5$)	4.94	0.026	0.074	−0.29	−0.49	−1.20 (0.48)	+0.25	−1.45
Constant G	4.94	0.000	0.055	—	−0.59	−3.11 (0.42)	−3.63	+0.53
Panel B: Smooth τ_ℓ, G residual								
Output-indexed τ_ℓ ($\eta_\tau = 0.5$)	4.94	0.076	0.042	+0.26	−0.37	−0.47 (0.21)	+0.27	−0.74
Constant τ_ℓ	4.94	0.056	0.033	+0.25	−0.60	−1.59 (0.22)	−0.76	−0.84

Notes: Headline CEVs are the cross-seed mean (standard deviation) of the ergodic-mean CEV on the 60×48 grid under the common protocol ($T = 10{,}000$, burn-in 500, household-resolve, mean G/Y one-pass scaling at seed 4242; $n = 20$ seeds). The λ_C/λ_G columns are the exact two-channel split of each CEV: with separable log utility and common exogenous paths, $(1+\lambda) = (1+\lambda_C)(1+\lambda_G)$, where λ_C prices the counterfactual's private-consumption path and λ_G its public-consumption path (both in private-consumption units); the split is exact seed by seed. Moment columns are cross-seed means over the same 20 seeds as the CEVs, from the same household-resolve simulation pipeline. Conventions for the moment columns: $\sigma(\tilde{Y})$ and $\text{corr}(G,Y)$ use HP($\lambda = 100$) cyclical components of the log series; $\sigma(G)$ and $\sigma(\tau_\ell)$ are level standard deviations; $\text{corr}(\tau_\ell, Y)$ is computed between the *level* of τ_ℓ and the HP cycle of $\log Y$. In Panels A and B the reported τ_ℓ is the *effective implementing tax* recovered ex post from the period budget identity (15) given the realized path of G (including the adjustment-cost term, at date-t wages, with G_t predetermined by the date-$(t-1)$ choice); it is *not* the rule target (see the text paragraph "Why the Panel B tax columns differ from the rule targets"). Source: replication package (financing-assignment moment table, 60×48). Appendix I reports the same four cells under the alternative asset normalization, which leaves them essentially unchanged, and under a re-allocation of the same number of purchases nodes; Section 6.4 explains why the ergodic cells are quoted as a band rather than as point estimates.

The constant-purchases cell is a level, not a cycle. Under the one-pass protocol the constant level chosen at the anchor seed realizes a mean G/Y of 0.328 against the benchmark's 0.313, so the rule delivers slightly more public consumption than the equilibrium and finances it with a mean labor tax of 0.501 against 0.479. That is the whole of the cell: the positive λ_G is the extra public consumption, and the large negative λ_C is the tax that pays for it. Iterating the level to a fixed point, so the rule realizes the benchmark share rather than targeting it, moves the cell from −3.11% to −1.86% and brings it inside the band occupied by the other three.

The sensitivity to the grid is larger than that adjustment. The fixed-point level, 0.506, falls between two adjacent purchases nodes. Evaluated at the lower node, $G = 0.455$, the constant-purchases rule is worth +5.01%; at the upper node, $G = 0.533$, it is worth −3.31%. One grid step in the level, 0.078 in G or about five percent of steady-state output, swings the comparison by 8.3 percentage points, whereas the entire cyclical content of the four-rule table spans about two. We therefore report comparisons involving constant G with an explicit level qualification and do not read its rank in Table 6 as informative about the value of smoothing purchases. The same arithmetic is the first sign of the result developed in Section 6.4: in this economy the level of public provision dominates its cyclical profile by an order of magnitude.

As a robustness check, we add the tax-deviation term (26) with $\lambda_\tau = 0.50$ and re-solve the Markov-perfect equilibrium from scratch, with everything else unchanged. On the full 60×48 state space, the resulting purchases policy reproduces the no-penalty baseline exactly (maximum absolute difference zero), and the measured welfare difference between the two benchmark allocations is zero

on all 20 evaluation seeds. Independent initializations produce the same policy, so this invariance is not a warm-start artifact. The term is therefore inactive at the margin the government actually chooses on: the implementing tax is pinned by the budget identity given purchases, so penalizing its deviation does not change which purchases node is optimal. Every welfare comparison in the paper is accordingly unchanged by its inclusion.

The added term is not irrelevant at every point. Recomputing the government's one-step discrete best response at the converged value function, with the argmax-stabilization devices disabled, shows that it changes the argmax in 4.5% of state cells (1.4% when weighted by ergodic visitation), by an average of 2.6 grid steps conditional on a change. These local changes are absorbed by the equilibrium policy approximation and leave the converged policy unchanged. Within the numerical protocol, adding the reduced-form tax-deviation term therefore does not change the headline equilibrium or its welfare comparison.

This conclusion applies to the ergodic comparison. The measure excludes the transition into each rule's long-run distribution; that omission has no determinate sign, although it favors a reform if adoption is costly. The policy claim in the abstract concerns the countercyclical reform that corresponds to the standard emerging-market "smoothing G" counterfactual. Attributing its entire cost to procyclicality would miss the assignment result, which is about incidence rather than sign: comparable cyclical adjustment on the financing side costs about as much on average and imposes it in the opposite states.

Why the Panel B tax columns differ from the rule targets. A reader who checks Table 6 against the rule definitions will notice that the B.i row reports $\mathrm{corr}(\tau_\ell, Y) = -0.37$ although the B.i target $\tau_\ell^B \, (Y_t/Y_{\mathrm{ref}})^{0.5}$ satisfies $\mathrm{corr}(\text{target}, Y) = +1$ by construction, and that the B.ii row reports $\sigma(\tau_\ell) = 0.033$ although the B.ii target is a constant. The rules are implemented exactly as stated: the tax target determines residual purchases G_{t+1} from date-t revenue. The τ_ℓ reported in the moment columns, however, is the *effective implementing tax*, recovered ex post from the period budget identity (15) given the realized path of G: it finances the predetermined G_t (chosen from date-$(t-1)$ revenue under the rule) plus the realized adjustment cost, at date-t wages. This is the same convention used for the baseline row and Panel A, which keeps the columns comparable across all five rows. Three mechanical wedges separate the effective tax from the rule target: the one-period timing shift (G_t embeds the date-$(t-1)$ target), the date-t wage in the denominator of (15) (a boom raises w_t and mechanically lowers the implementing rate for a given G_t), and the clipping of residual G at zero and at the grid bounds. The wage-denominator effect dominates, which is why $\mathrm{corr}(\tau_\ell, Y)$ is negative even under a target whose correlation with output is $+1$, and why $\sigma(\tau_\ell) > 0$ under a constant target. One potential further wedge is absent: the $\bar{\tau}_\ell^{\max} = 0.50$ cap on the rule *target* never binds along the simulated paths (the target reaches 0.499 in 0.00% of periods across all 20 seeds), so the Panel B rules operate in their interior everywhere and the cap does not interact with the anchor-drift discussion in Appendix I. The welfare numbers are unaffected by this reporting convention: in the Panel B household re-solves the household faces the

rule-consistent tax process, and the CEVs compare utility paths, not the moment columns.

Where the asymmetry bites: state-conditional welfare. The ergodic averages in Table 6 hide differences across states that are an order of magnitude larger than the averages themselves, and it is these differences, rather than the averages, that carry the policy content. The two assignments cost about the same on average. They do not cost the same in the same states. The state-conditional decomposition below shows that the spending-side rule concentrates its financing squeeze in low-productivity, high-world-rate states, where the household's marginal utility of consumption is highest, whereas the financing-side rule leaves those states nearly untouched and concentrates its cost in high-productivity states instead. The assignment is therefore a choice about the timing of the burden across the cycle, not about its size.

Figure 3 shows where the welfare cost is concentrated. We group the simulated observations by their exogenous node indices. "Low z" combines the bottom two of the three TFP nodes, at or below the median, while "high z" is the top node. Low and high r^* refer to the two world-rate nodes. In the order (low z, low r^*), (low z, high r^*), (high z, low r^*), and (high z, high r^*), the four groups cover 69.2%, 5.6%, 23.4%, and 1.9% of post-burn observations. The high-r^* groups are small because the world-rate process places stationary probability 0.075 on the expensive node (Section 4).

Following the welfare protocol above, we calculate (29) within each group on the headline 60×48 grid using the same 20 seeds as Table 6. The exact aggregation identity is in log-CEV units, not CEV levels: the $\overline{\theta}$-weighted average of the conditional values of $\log(1+\lambda)$ equals the unconditional $\log(1+\lambda)$. Applying the exponential transformation then recovers the unconditional CEV, with numerical identity error below 10^{-16}.

The spending-side rule's loss occurs almost entirely in one group. Its conditional CEV is -12.6% in the low-z, high-r^* expensive-financing episodes, against -0.4% in the low-z, low-r^* group that covers most of the ergodic distribution, -1.6% in the high-z, low-r^* group, and $+1.0\%$ in the small high-z, high-r^* group. The financing-side rule inverts this profile almost exactly. It gains $+4.3\%$ in the low-z, low-r^* group, is close to neutral at -0.3% in the expensive-financing episodes where the spending-side rule loses most, and loses -13.0% and -6.1% in the two high-productivity groups. Each of these conditional values is large relative to its cross-seed standard deviation, which runs from 0.35 to 1.03 percentage points, so the contrast between the two profiles is identified even though the ergodic averages they aggregate to are not separated at that precision.

Within that low-z, high-r^* group, the spending-side loss comes almost entirely from the private-consumption channel ($\lambda_C = -12.0$ of a total of -12.6). The rule keeps purchases high when revenue is weak and external financing is expensive, so the residual tax rises precisely when the household's marginal utility of consumption is highest. The unconditional $\lambda_C = +0.25$ therefore understates what the rule does to households: the financing burden is not passed back on average, but it is passed back sharply in the episodes that matter most. The financing-side rule's losses have the same private-consumption source in the states where they occur ($\lambda_C = -17.4$ in the high-z, low-r^*

group) but there the marginal utility of consumption is low, which is why the same mechanism produces a smaller welfare cost on average.

The main point of the figure is that welfare effects are distributed unevenly across exogenous states, and in opposite directions under the two assignments. The spending-side rule imposes its loss where fiscal constraints are tightest; the financing-side rule imposes its loss where they are loosest. For an economy exposed to global financial fragmentation, which raises the frequency and severity of exactly the low-z, high-r^* episodes, that difference is the substance of the assignment choice even though the two rules are nearly indistinguishable on an ergodic average.

Implications for the policy reading. A reform that smooths public purchases over the cycle shifts the residual financing adjustment onto the labor tax. It does not necessarily increase the tax's unconditional volatility, but it changes the states in which the financing burden is imposed. Welfare must therefore be evaluated for the joint (G, τ_ℓ) package.

Under the headline household-resolve protocol the spending-side reform has a CEV of -1.2% and its financing-side counterpart -0.5%. Both lower welfare on the headline grid; the half-point that separates them lies within the discretization uncertainty of Appendix C. For fiscal-rule design in developing economies exposed to world-rate risk the two are nonetheless not interchangeable, but the reason is distributional rather than a difference in average cost: the spending-side reform withdraws household resources in the expensive-financing episodes, and the financing-side reform withdraws them in booms. A policymaker who weights bad states more heavily than the household's own utility function does—which is the usual motivation for a fiscal rule in the first place—will not be indifferent between two reforms that cost the same on average.

World-rate fragmentation (qualitative). If expensive world-rate episodes become more severe or frequent, fiscal headroom becomes more valuable and feasibility in low-output states is reduced. In our model, these episodes are rare and independent (Section 4), so we focus on how their severity and frequency affect outcomes, not their serial dependence. When world-rate fragmentation increases, how cyclical adjustment is split between purchases and financing matters as much as the cyclical stance of G alone.

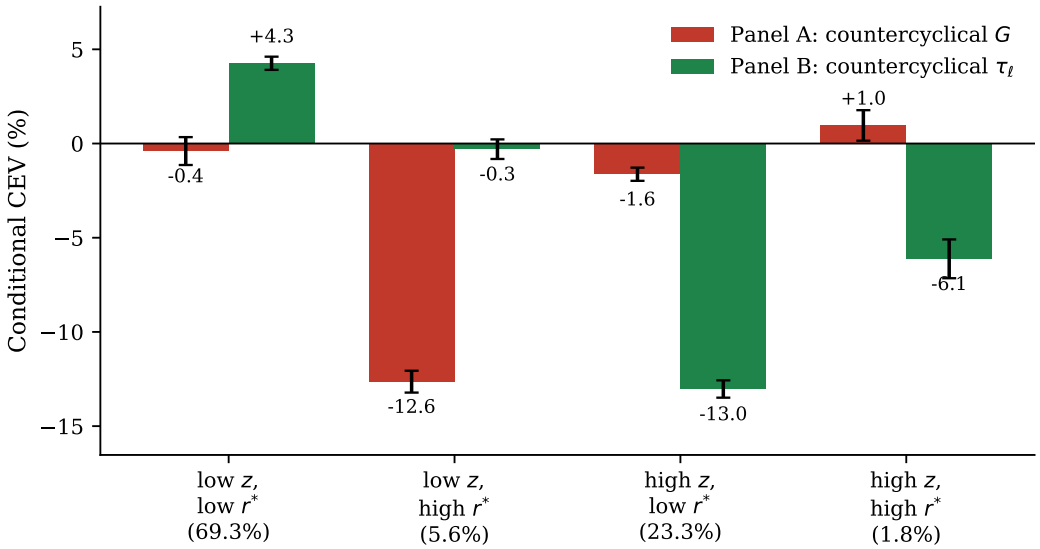


Figure 3: State-conditional CEV (%) of the Panel A and Panel B rules against the time-consistent benchmark, conditional on the exogenous node bins described in the text (low z pools the bottom two TFP nodes; the r^* bins are the two world-rate nodes); headline 60×48 grid, unified pipeline, cross-seed means over the same 20 seeds as Table 6, with cross-seed standard-deviation error bars. Bin sample shares are printed under each label. Conditional CEV levels do not add linearly; the $\overline{\theta}$-weighted conditional log CEVs aggregate exactly to the unconditional log CEV, which is then transformed back into the unconditional CEV in Table 6. Panel A (countercyclical G, red) concentrates its loss in the (low z, high r^*) bin (-12.6%, cross-seed sd about 0.6pp), the rare expensive-financing episodes in which fiscal feasibility tightens hardest, and is close to neutral elsewhere. Panel B (output-indexed tax rule, green) has the mirror-image profile: it is near-neutral in the expensive-financing bin and loses in the high-productivity bins. The two high-r^* cells are computed from smaller sub-samples; their cross-seed standard deviations are correspondingly larger (about 0.5–1.0pp for Panel B).

6.4 The cost of discretion is a level, not a cycle

Every comparison so far holds the long-run mean of G/Y fixed, which is what makes them comparisons of the assignment. That protocol is also what hides the largest number in the paper. This subsection removes the restriction and asks what the *level* of public provision costs under discretion, and why a government that would gain from changing it does not. We do not compute a full-commitment planner; the objects below are a welfare-maximizing spending share computed from the equilibrium's own resources, permanent proportional restraint of the discretionary rule, and the value of commitment implied by comparing sustained restraint to a one-period deviation.[9]

Welfare-maximizing spending share. At each date and state we take total absorption $C_t + G_t$ as realized in the benchmark simulation and re-split it between private and public consumption to maximize the period utility $\theta_t \log C + (1 - \theta_t) \log G$. The re-split is feasible period by period with resources the equilibrium itself generates, so no reassignment between purchases and the household can do better within that resource path. Evaluated with the same transform as the rule cells, this welfare-maximizing spending share is worth +21.9% of consumption relative to discretion (cross-seed standard deviation 0.36), and it moves mean G/Y from 0.313 to 0.168. Allowing total absorption to be fully smoothed as well changes the gain only to +20.5%, so essentially all of the gain is the intratemporal split and almost none of it is intertemporal smoothing.

Permanent proportional restraint. The welfare-maximizing spending share optimizes freely each period and is not itself a policy. A restricted alternative is to scale the equilibrium's own state-contingent purchases rule by a constant factor κ, which preserves its cyclical shape exactly and moves only its level—permanent proportional restraint. We re-solve the household under each scaled rule and simulate under the same protocol. Table 7 reports the result. The gain is hump-shaped and peaks at +13.9% of consumption at $\kappa = 0.55$, where mean G/Y is 0.172. That the peak of a dynamic experiment with the household responding falls within four thousandths of the independently computed welfare-maximizing share of 0.168, given that the two computations share no code, is a useful check on both.

Value of commitment. The last column of Table 7 prices the same restraint when it lasts one period. At $\kappa = 0.55$ a government that cuts purchases today and then lets the equilibrium policy resume gives up 0.47% of consumption by its own objective; sustaining the same cut pays 13.9%. The ratio is about thirty to one—the value of commitment measured in the model's own units. The deviation is nearly worthless because next period's government inherits the state and reverts, so the deviator buys a single period of a different split and bears the transition. This is the time-consistency trap in its plainest form, and it is the reason the level distortion survives in an

[9] A full-commitment planner with incomplete markets and constant debt did not converge on our state space at a resolution comparable to the equilibrium; solving that problem is left for future work. The permanent-proportional-restraint exercise is the implementable bound we use instead.

Table 7: Permanent proportional restraint and the one-period deviation

		Sustained (CEV, %)			One period (%)
κ	mean G/Y	total	private	public	median
1.00	0.312	—	—	—	—
0.90	0.281	+4.99	+8.72	−3.43	−0.10
0.80	0.250	+9.00	+17.37	−7.13	−0.15
0.70	0.219	+11.94	+25.99	−11.15	−0.23
0.60	0.187	+13.63	+34.60	−15.58	−0.32
0.55	0.172	**+13.92**	+38.90	−17.98	−0.47
0.50	0.156	+13.79	+43.19	−20.53	−0.63

Notes. Headline 60×48 grid with the asset reference point $A_{\text{ref}} = K_{\text{ref}} + \bar{B}$, ten seeds, common protocol. "Sustained" is permanent proportional restraint: scale the equilibrium purchases rule by κ at every date and state, with the household re-solved under the scaled rule; cross-seed standard deviations rise from 0.05 at $\kappa = 0.90$ to 0.40 at $\kappa = 0.50$. "One period" evaluates the government's own objective at the node nearest κ times its chosen purchases, with the equilibrium policy resuming immediately afterward, and is reported as the median over 8,664 on-path state cells; the distribution is left-skewed, with a mean of −2.3% driven by states near the feasibility boundary. The ratio of sustained to one-period values at $\kappa = 0.55$ is the value of commitment. Both columns use the transform of (29).

equilibrium whose government is fully benevolent and faces no political-economy or default friction at all.

Two implications follow. The first is substantive: in this class of models the cost of limited commitment is overwhelmingly a level phenomenon, and cyclical-rule reform of the kind the policy debate concentrates on operates on the smaller margin by roughly an order of magnitude. The second is methodological, and it resolves what would otherwise look like a numerical failure. The government's objective is nearly flat in its choice of purchases—the median on-path separation between the best node and the runner-up is worth 0.10% of consumption on the headline grid (0.02% after the node re-allocation; Appendix C)—and that flatness is not an artifact of the discretization. It is the same fact as the last column of Table 7, seen from a different angle. An objective whose one-period deviations are nearly worthless is an objective that is nearly flat. The discretization inherits the flatness rather than causing it, which is why refining the grid multiplies near-ties instead of resolving them, and why we report the ergodic rule cells as bounded.

7 Conclusion

A calibrated time-consistent small-open-economy model in which a Markov-perfect fiscal authority chooses government purchases under fundamental shocks generates procyclical purchases as the optimal equilibrium response. Without political-economy distortions or sovereign-default risk, relative volatility of public purchases is 1.94 against 2.20 in annual emerging-market data (Section 6.2). Procyclical fiscal policy in emerging markets is therefore consistent with optimization amid limited

commitment even in an environment with no political-economy or default frictions; the channels emphasized in the existing literature are not the only mechanisms that generate it.

The paper's central quantitative finding is that the cost of limited commitment in this economy is a level rather than a cycle. The discretionary government provides public goods at 31% of output, while the welfare-maximizing spending share puts them at 17%. Permanent proportional restraint of the equilibrium's purchases rule down to that level, holding its cyclical shape fixed, is worth +13.9% of consumption. No government choosing one period at a time will do it: the same restraint sustained for a single period is worth −0.47% by the government's own objective, about one thirtieth of what sustaining it pays, because the next period's government inherits the state and reverts. That ratio is the value of commitment measured in the model's own units, and it obtains in an economy whose government is fully benevolent and faces neither political-economy nor default frictions. It also explains, rather than excuses, the numerical flatness that limits the precision of the cyclical comparisons: an objective whose one-period deviations are nearly worthless is a flat objective.

Against that benchmark, the cyclical reassignments the policy debate concentrates on operate on the smaller margin. The four ergodic cells are not identified to a sign, let alone a ranking: they are all negative on the headline grid, and three of the four change sign when the same nodes are re-allocated over the ergodic range of G (Table 12). Neither side of the budget is the cheap side on which to place the cyclical adjustment.

What the assignment does determine, sharply and with tight standard errors, is which states bear the cost. The spending-side rule concentrates its loss in the low-productivity, high-world-rate states in which fiscal feasibility tightens, at −12.6% of consumption, and falls there on the household's private-consumption path. The financing-side rule leaves those states nearly untouched at −0.3% and instead loses −13.0% in high-productivity states, gaining +4.3% in the low-rate states that make up most of the ergodic distribution. Two reforms that are indistinguishable on an ergodic average are therefore very different reforms, and a policymaker who cares about outcomes in bad states will not be indifferent between them.

For developing economies that face higher and more variable world real interest rates, the welfare implications are direct. The in-sample evidence is the state-conditional incidence just described: the spending-side cost is concentrated in precisely the expensive-financing states, so economies in which those states are more severe or more frequent draw more heavily on the region where the reform is costly and the financing-side alternative least so. We state the corresponding comparative static qualitatively, as a direction the conditional decomposition supports rather than a magnitude the paper estimates. Fiscal-rule design in such settings should be framed as a choice between financing-side smoothing (of the labor tax or of debt) and spending-side smoothing (of public purchases), with the cyclical stance of one pinning down the cyclical stance of the other. Reforms that smooth purchases against the cycle without redesigning how that smoothing is financed deliver the destabilizing financing path that the welfare comparison isolates, most aggressively in exactly the states where world-rate fragmentation tightens fiscal feasibility.

The quantitative results rest on the modeling choices in Section 4: the constant-debt rule as the financing closure; the inelastic-labor specification under which the implementing labor tax is a residual financing wedge rather than a labor-supply tax distortion; the absence of sovereign default, of an active monetary block, and of a debt-elastic interest-rate premium; and the reliance on low-dimensional parametric summaries of the policy rules. Because the economy is stationary through impatience rather than through a premium that anchors the external position, consumption-equivalent levels are evaluated at a higher marginal utility than the data imply (Section I.1); the ranking of level against cycle, and the opposite state-conditional incidence within this economy, do not depend on that scaling. A full-commitment planner over purchases and the labor tax jointly is left for future work: that problem did not converge on our state space at a resolution comparable to the equilibrium, and the permanent-proportional-restraint exercise of Section 6.4 is the implementable bound we use instead. An elastic-labor extension would add a labor-supply distortion margin on top of the financing-wedge channel identified here. A state-contingent debt-issuance closure would allow debt to absorb part of the cycle and reweight the spending–financing–debt margins; the analytical Remark in Section 4.6 shows the financing-assignment channel survives any debt-issuance regime short of the perfectly state-contingent response, the single-shock analogue of the Lucas–Stokey complete-markets benchmark.

References

Philippe Aghion, Ioana Marinescu, Ricardo J Caballero, and Anil K Kashyap. Cyclical budgetary policy and economic growth: What do we learn from oecd panel data?[with comments and discussion]. *NBER Macroeconomics annual*, 22:251–297, 2007.

Mark Aguiar and Gita Gopinath. Emerging market business cycles: The cycle is the trend. *Journal of political Economy*, 115(1):69–102, 2007.

S Rao Aiyagari, Albert Marcet, Thomas J Sargent, and Juha Seppälä. Optimal taxation without state-contingent debt. *Journal of Political Economy*, 110(6):1220–1254, 2002.

Enrique Alberola, Yavuz Arslan, Gong Cheng, and Richhild Moessner. Fiscal response to the covid-19 crisis in advanced and emerging market economies. *Pacific Economic Review*, 26(4):459–468, 2021.

Alberto Alesina, Filipe R Campante, and Guido Tabellini. Why is fiscal policy often procyclical? *Journal of the european economic association*, 6(5):1006–1036, 2008.

James E Alt and David Dreyer Lassen. Transparency, political polarization, and political budget cycles in oecd countries. *American Journal of Political Science*, 50(3):530–550, 2006.

Martín Ardanaz, Eduardo Cavallo, Alejandro Izquierdo, and Jorge Puig. Growth-friendly fiscal rules? safeguarding public investment from budget cuts. *IDB Working Paper*, 2020. URL `https://www.bde.es/f/webpi/SES/seminars/2020/Fich/sie201021.pdf`.

Cristina Arellano. Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712, 2008.

Rüdiger Bachmann and Jinhui H Bai. Politico-economic inequality and the comovement of government purchases. *Review of Economic Dynamics*, 16(4):565–580, 2013a.

Rüdiger Bachmann and Jinhui H Bai. Public consumption over the business cycle. *Quantitative Economics*, 4(3):417–451, 2013b.

Robert J Barro. On the determination of the public debt. *Journal of political Economy*, 87(5, Part 1):940–971, 1979.

Robert B Barsky and Eric R Sims. News shocks and business cycles. *Journal of monetary Economics*, 58(3):273–289, 2011.

Marianne Baxter and Robert G King. Fiscal policy in general equilibrium. *The American economic review*, pages 315–334, 1993.

Paul Beaudry and Franck Portier. Stock prices, news, and economic fluctuations. *American economic review*, 96(4):1293–1307, 2006.

Paul Beaudry, Martial Dupaigne, and Franck Portier. Modeling news-driven international business cycles. *Review of Economic Dynamics*, 14(1):72–91, 2011.

Javier Bianchi, Pablo Ottonello, and Ignacio Presno. Fiscal stimulus under sovereign risk. *Journal of Political Economy*, 131(9):2328–2369, 2023.

Henning Bohn. The behavior of u.s. public debt and deficits. *Quarterly Journal of Economics*, 113 (3):949–963, 1998. doi: 10.1162/003355398555793.

Eleonora Bova, Nathalie Carcenac, and Martine Guerguil. Fiscal rules and the procyclicality of fiscal policy in the developing world. IMF Working Paper 14/122, International Monetary Fund, 2014.

Guillermo A Calvo, Alejandro Izquierdo, and Luis-Fernando Mejia. On the empirics of sudden stops: the relevance of balance-sheet effects, 2004.

Francesca Caselli and Julien Reynaud. Do fiscal rules cause better fiscal balances? a new instrumental variable strategy. *European Journal of Political Economy*, 63:101873, 2020. doi: 10.1016/j.ejpoleco.2020.101873.

Serhan Cevik. Anchor me: the benefits and challenges of fiscal responsibility. *Asian-Pacific Economic Literature*, 33(1):33–47, 2019.

Giancarlo Corsetti and Gernot J Müller. Multilateral economic cooperation and the international transmission of fiscal policy. In *Globalization in an age of crisis: Multilateral economic cooperation in the twenty-first century*, pages 257–297. University of Chicago Press, 2013.

Gabriel Cuadra, Juan M Sanchez, and Horacio Sapriza. Fiscal policy and default risk in emerging markets. *Review of Economic Dynamics*, 13(2):452–469, 2010.

Davide Debortoli and Ricardo Nunes. Fiscal policy under loose commitment. *Journal of Economic Theory*, 145(3):1005–1032, 2010. doi: 10.1016/j.jet.2009.12.004.

Margarita Debuque-Gonzales, C. J. Diokno-Sicat, J. P. Corpus, R. H. Palomar, M. G. Ruiz, and R. M. Miral. Fiscal effects of the COVID-19 pandemic: Assessing public debt sustainability in the Philippines. Research Paper Series 2026-01, Philippine Institute for Development Studies, Quezon City, 2026. Earlier version: PIDS Discussion Paper Series No. 2022-17, `https://doi.org/10.62986/dp2022.17`.

Department of Budget and Management. Budget of expenditures and sources of financing (besf), various fiscal years. `https://www.dbm.gov.ph/`, 2025. URL `https://www.dbm.gov.ph/`. Annual budget documentation providing functional and economic spending classifications; accessed September 4, 2025.

Antonio Fatás and Ilian Mihov. Policy volatility, institutions, and economic growth. *Review of Economics and Statistics*, 95(2):362–376, 2013.

Jesús Fernández-Villaverde, Pablo Guerrón-Quintana, Keith Kuester, and Juan Rubio-Ramírez. Fiscal volatility shocks and economic activity. *American Economic Review*, 105(11):3352–3384, 2015.

Jeffrey A Frankel. A solution to fiscal procyclicality: The structural budget institutions pioneered by chile. Technical report, National Bureau of Economic Research, 2011.

Jeffrey A Frankel, Carlos A Vegh, and Guillermo Vuletin. On graduation from fiscal procyclicality. *Journal of Development Economics*, 100(1):32–47, 2013.

Michael Gavin and Roberto Perotti. Fiscal policy in latin america. *NBER macroeconomics annual*, 12:11–61, 1997.

Atish R. Ghosh, Jun I. Kim, Enrique G. Mendoza, Jonathan D. Ostry, and Mahvash S. Qureshi. Fiscal fatigue, fiscal space and debt sustainability in advanced economies. *The Economic Journal*, 123(566):F4–F30, 2013. doi: 10.1111/j.1468-0297.2012.02553.x.

Douglas Gollin. Getting income shares right. *Journal of political Economy*, 110(2):458–474, 2002.

Martine Guerguil, Pierre Mandon, and René Tapsoba. Flexible fiscal rules and countercyclical fiscal policy. *Journal of Macroeconomics*, 52:189–220, 2017.

Friedrich Heinemann, Marc-Daniel Moessinger, and Mustafa Yeter. Do fiscal rules constrain fiscal policy? a meta-regression analysis. *European Journal of Political Economy*, 51:69–92, 2018. doi: 10.1016/j.ejpoleco.2017.03.008.

Fédéric Holm-Hadulla, Sebastian Hauptmeier, and Philipp Rother. The impact of expenditure rules on budgetary discipline over the cycle. *Applied Economics*, 44(25):3287–3296, 2012. doi: 10.1080/00036846.2011.572855.

Ethan Ilzetzki and Carlos A Végh. Procyclical fiscal policy in developing countries: Truth or fiction? Technical report, National Bureau of Economic Research, 2008.

Ethan Ilzetzki, Enrique G Mendoza, and Carlos A Végh. How big (small?) are fiscal multipliers? *Journal of monetary economics*, 60(2):239–254, 2013.

International Monetary Fund. Government finance statistics manual 2014. Technical report, International Monetary Fund, 2014. GFSM 2014 statistical concepts and classifications.

International Monetary Fund. Philippines: Public investment management assessment (pima), 2019. URL `https://www.imf.org/en/Publications/WP/Issues/2019/03/25`.

International Monetary Fund Independent Evaluation Office. Evaluation of Prolonged Use of IMF Resources. Independent evaluation office report, Independent Evaluation Office, International Monetary Fund, Washington, DC, sep 2002. URL `https://www.imf.org/external/pubs/ft/epui/2002/pdf/report.pdf`. Includes Philippines case study (23 arrangements, 1962–2000).

Dany Jaimovich and Ugo Panizza. Procyclicality or reverse causality? Technical report, Working paper, 2007.

Nir Jaimovich and Sergio Rebelo. Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118, 2009.

João Tovar Jalles, Youssouf Kiendrebeogo, Raphael W. Lam, and Roberto Piazza. Revisiting the countercyclicality of fiscal policy. *Empirical Economics*, 67:877–914, 2024. doi: 10.1007/s00181-024-02586-z.

Kenneth L Judd, Lilia Maliar, and Serguei Maliar. Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models. *Quantitative Economics*, 2(2):173–210, 2011.

Güneş Kamber, Konstantinos Theodoridis, and Christoph Thoenissen. News-driven business cycles in small open economies. *Journal of International Economics*, 105:77–89, 2017.

Graciela L. Kaminsky, Carmen M. Reinhart, and Carlos A. Végh. When it rains, it pours: Procyclical capital flows and macroeconomic policies. In Mark Gertler and Kenneth Rogoff, editors, *NBER Macroeconomics Annual 2004, Volume 19*, volume 19, pages 11–82. MIT Press, 2005.

Paul Klein, Per Krusell, and Jose-Victor Rios-Rull. Time-consistent public policy. *The Review of Economic Studies*, 75(3):789–808, 2008.

Gary Koop, M. Hashem Pesaran, and Simon M. Potter. Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1):119–147, 1996.

Karen A Kopecky and Richard MH Suen. Finite state markov-chain approximations to highly persistent processes. *Review of Economic Dynamics*, 13(3):701–714, 2010.

Philip R Lane. The cyclical behaviour of fiscal policy: evidence from the oecd. *Journal of Public economics*, 87(12):2661–2675, 2003.

Philip R. Lane and Gian Maria Milesi-Ferretti. The external wealth of nations revisited: International financial integration in the aftermath of the global financial crisis. *IMF Economic Review*, 66(1):189–222, 2018. doi: 10.1057/s41308-017-0048-y.

Eric M. Leeper. Equilibria under "active" and "passive" monetary and fiscal policies. *Journal of Monetary Economics*, 27(1):129–147, 1991. doi: 10.1016/0304-3932(91)90007-B.

Eric M. Leeper, Todd B. Walker, and Shu-Chun Susan Yang. Government investment and fiscal stimulus. *Journal of Monetary Economics*, 57(8):1000–1012, 2010. doi: 10.1016/j.jmoneco.2010.09.002.

Eric M. Leeper, Todd B. Walker, and Shu-Chun Susan Yang. Fiscal foresight and information flows. *Econometrica*, 81(3):1115–1145, 2013. doi: 10.3982/ECTA8337.

Norman V. Loayza and Claudio Raddatz. The structural determinants of external vulnerability. *World Bank Economic Review*, 21(3):359–387, 2007.

Robert E. Lucas and Nancy L. Stokey. Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics*, 12(1):55–93, 1983.

Francisco Arroyo Marioli, Antonio Fatas, and Garima Vasishtha. Fiscal policy volatility and growth in emerging markets and developing economies. *International Review of Economics & Finance*, 92:758–777, 2024.

Eric Maskin and Jean Tirole. Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory*, 100(2):191–219, October 2001. doi: 10.1006/jeth.2000.2785.

Enrique G. Mendoza. Real business cycles in a small open economy. *The American Economic Review*, 81(4):797–818, 1991.

Enrique G. Mendoza and P. Marcelo Oviedo. Fiscal policy and macroeconomic uncertainty in developing countries: The tale of the tormented insurer. *NBER Working Paper*, 12586(12586), 2006.

Johanness Mueller, Renaud Duplay, Luc Eyraud, Jason Harris, Murray Petrie, and Sage de Clerck. Philippines: Fiscal transparency evaluation. IMF Country Report 15/156, International Monetary Fund, 2015.

Pablo A. Neumeyer and Fabrizio Perri. Business cycles in emerging economies: The role of interest rates. *Journal of Monetary Economics*, 52(2):345–380, 2005.

Maurice Obstfeld and Kenneth S. Rogoff. *Foundations of International Macroeconomics*. MIT Press, Cambridge, MA, 1996. ISBN 9780262150477.

Stephane Pallage and Michel A Robe. On the welfare cost of economic fluctuations in developing countries. *International Economic Review*, 44(2):677–698, 2003.

Philippine Statistics Authority. Philippine system of national accounts: Sources and methods. Technical report, Philippine Statistics Authority, Quezon City, Philippines, 2020.

Demian Pouzo and Ignacio Presno. Optimal taxation with endogenous default under incomplete markets. *American Economic Journal: Macroeconomics*, 14(3):1–41, 2022.

Alvaro Riascos and Carlos A Vegh. Procyclical government spending in developing countries: The role of capital market imperfections. *unpublished (Washington: International Monetary Fund)*, 2003.

K Geert Rouwenhorst. Asset pricing implications of equilibrium business cycle models. *Frontiers of business cycle research*, 1:294–330, 1995.

Ms Andrea Schaechter, Mr Tidiane Kinda, Mrs Nina Budina, and Anke Weber. *Fiscal Rules in Response to the Crisis: Toward the" Next-Generation" Rules: A New Dataset*. International Monetary Fund, 2012.

Stephanie Schmitt-Grohé and Martín Uribe. Closing small open economy models. *Journal of International Economics*, 61(1):163–185, 2003. doi: 10.1016/S0022-1996(02)00056-9.

Stephanie Schmitt-Grohé and Martin Uribe. What's news in business cycles. *Econometrica*, 80(6): 2733–2764, 2012.

Ernesto Talvi and Carlos A. Végh. Tax base variability and procyclical fiscal policy in developing countries. *Journal of Development economics*, 78(1):156–190, 2005.

George Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters*, 20(2):177–181, 1986.

Aaron Tornell and Philip R. Lane. The voracity effect. *American Economic Review*, 89(1):22–46, 1999.

Martín Uribe and Vivian Z. Yue. Country spreads and emerging countries: Who drives whom? *Journal of International Economics*, 69(1):6–36, 2006.

U.S. Bureau of Economic Analysis. Government consumption expenditures and gross investment. In *NIPA Handbook: Concepts and Methods of the U.S. National Income and Product Accounts*, chapter 9. U.S. Bureau of Economic Analysis, 2024. URL `https://www.bea.gov/resources/methodologies/nipa-handbook/pdf/chapter-09.pdf`. Updated December 2024.

Carlos A Vegh and Guillermo Vuletin. How is tax policy conducted over the business cycle? *American Economic Journal: Economic Policy*, 7(3):327–370, 2015.

Iván Werning. Optimal fiscal policy with redistribution. *The Quarterly Journal of Economics*, 122 (3):925–967, 2007.

Jaejoon Woo. Why do more polarized countries run more procyclical fiscal policy? *The Review of Economics and Statistics*, 91(4):850–870, 2009.

World Bank. World development indicators (wdi): Gdp per capita (ny.gdp.pcap.cd); trade (% of gdp) (ne.trd.gnfs.zs); personal remittances, received (% of gdp) (bx.trf.pwkr.dt.gd.zs). `https://data.worldbank.org/`, 2025a. Accessed September 4, 2025.

World Bank. Worldwide governance indicators (wgi): Country data and metadata. `https://info.worldbank.org/governance/wgi/`, 2025b. Accessed September 4, 2025.

A Proof of Lemma 1

Proof sketch. Start by taking the logarithm of the period budget identity, $\tau_{\ell,t}\, w_t L = G_t + AC_t + r_t^* \bar{B} - \tau_a r_t^* A_t$. Then, differentiate $\log(\tau_{\ell,t}\, w_t L) = \log(G_t + AC_t + r_t^* \bar{B} - \tau_a r_t^* A_t)$ with respect to $\log Y_t$ at the deterministic reference point. On the left side, we get $\varepsilon_\tau + \varepsilon_w$ because $\tau_\ell\, w\, L$ is a product. For the asset-income term, the first-order change is $\tau_a \big(r^*_{\text{ref}}\, dA_t + A_{\text{ref}}\, dr_t^*\big)$. The first part is zero because A_t is predetermined within the period, so $dA_t = 0$. The second part is also zero since $A_{\text{ref}} = 0$. Because $AC_{\text{ref}} = 0$ and dAC is a second-order change in dG, the right-hand side becomes

$$\frac{1}{G_{\text{ref}} + r^* \bar{B}} \left(G_{\text{ref}}\, \varepsilon_G \;+\; r^* \bar{B}\, \varepsilon_r\right) \;=\; \omega\, \varepsilon_G \;+\; (1-\omega)\, \varepsilon_r.$$

The second term comes from changes in debt payments, $r_t^* \bar{B}$, that vary with the cycle. Setting the left and right sides equal gives equation (27). The wage elasticity $\varepsilon_w = 1$ always holds, not just at the reference point. If we substitute the static capital demand $K_t = L(\alpha z_t/(r_t^* + \delta))^{1/(1-\alpha)}$ into the production function, we find $Y_t = z_t L\, (\alpha z_t/(r_t^* + \delta))^{\alpha/(1-\alpha)}$. The competitive wage from (10) is $w_t = (1-\alpha) z_t (\alpha z_t/(r_t^* + \delta))^{\alpha/(1-\alpha)}$, so $w_t = (1-\alpha)\, Y_t/L$ at every date and state. This means the wage is always proportional to output, no matter the shock, and $\varepsilon_w = d\log w_t / d\log Y_t = 1$ for any shock. □

B The deterministic reference point

At the deterministic reference point, the user cost sets the level of capital according to $F_K(K, 1) = r^* + \delta$. Using $\alpha = 0.33$, $\delta = 0.10$, and $\bar{r}^* = 0.04$, the capital–output ratio is $K_{\text{ref}}/Y_{\text{ref}} \approx 2.36$. When $\bar{B} = 0.5$ and the household is endowed with claims on capital and public debt, $A_{\text{ref}} = K_{\text{ref}} + \bar{B}$, the spending share $G_{\text{ref}}/Y_{\text{ref}} = 0.15$ is achieved with a labor tax of 21.9%. This reference rate, $\tau_{\ell,\text{ref}} = 0.219$, is also the center of the tax-deviation term used only in the robustness specification (26). Household-budget consumption and the resource-side residual coincide at $C_{\text{ref}} = 0.9371$ ($C/Y \approx 0.614$), because a zero external position implies balanced trade. The alternative $A_{\text{ref}} = 0$ assigns those claims to foreigners ($\text{NFA}_{\text{ref}} = -4.096$, or -268% of output) and implies a reference labor tax of 0.243; under that normalization the two consumption values differ by 0.164, or 10.7% of output. Table 8 shows the levels used as the normalization point for the numerical procedure.

Table 8: Quantities at the deterministic reference point

Variable	Value	Ratio / interpretation
Output Y_{ref}	1.5255	—
Capital K_{ref}	3.5958	$K/Y \approx 2.36$
Investment I_{ref}	0.3596	$I/Y \approx 0.236$
Government G_{ref}	0.2288	$G/Y = 0.15$
Consumption C_{ref} (household budget)	0.9371	$C/Y \approx 0.614$
Resource-side residual $Y_{\text{ref}} - I_{\text{ref}} - G_{\text{ref}}$	0.9371	Coincides with the budget value
Household assets $A_{\text{ref}} = K_{\text{ref}} + \bar{B}$	4.0958	Claims on capital and public debt
Net foreign assets $\text{NFA}_{\text{ref}} = A_{\text{ref}} - K_{\text{ref}} - \bar{B}$	0.0000	Balanced external position
Trade balance $TB_{\text{ref}} = -r^* \,\text{NFA}_{\text{ref}}$	0.0000	Balanced trade
Labor tax τ_ℓ	0.2194	21.9%

Notes: Values at the deterministic reference point implied by the calibration; this point is the ref normalization defined in Section 4 and constructed here. The household is endowed at the reference point with claims on the domestic capital stock and the outstanding public debt, $A_{\text{ref}} = K_{\text{ref}} + \bar{B}$, which is what sets the reference net foreign asset position to zero. The natural alternative, $A_{\text{ref}} = 0$, implies instead that foreigners own the entire capital stock and all public debt, giving a reference external position of -268% of output and a reference labor tax of 0.2434. Under the present normalization the household-budget value of consumption, $(1-\tau_{\ell,\text{ref}})\, w_{\text{ref}} L + (1-\tau_a)\, r^* A_{\text{ref}}$, and the resource-side residual $Y_{\text{ref}} - I_{\text{ref}} - G_{\text{ref}}$ coincide exactly, because a zero external position implies balanced trade; under $A_{\text{ref}} = 0$ they differ by 0.1638, or 10.7% of output. Appendix I reports the solved equilibrium under both normalizations and shows that the change leaves the policy, the moments and every welfare comparison intact. All quantities are reported in levels (output is normalized so that $Y_{\text{ref}} = 1.5255$). Ratios are computed from the reported values.

C Computational grids, state counts, and numerical accuracy

The basis transformation uses $\log A_t - a_{\text{norm}}$, where a_{norm} is the mean of $\log A_t$ for on-path observations with $A_t > 0$. If $A_t \leq 0$, we use $\log(\max\{A_t, 10^{-12}\})$ as a safeguard.

In the baseline simulation ($T = 10{,}000$, with a burn-in of 500), the ergodic distribution never reaches $A_t \leq 0$ (0% of periods). This means the basis is identified in the positive-A region where almost all the probability mass lies. We solve for and report both moments and welfare on the main 60×48 grid for (A, G). Appendix I checks robustness to the placement of those nodes.

The asset grid covers $[-1.5, 8.0]$, with the lower bound set by the borrowing limit (3). The government-purchases grid is constructed adaptively, with denser spacing near G_{ref} and bounds wide enough to include the simulated ergodic range of G. In the baseline, the simulated path for G spans $[0.229, 0.598]$ and never reaches a grid boundary. This means 0% boundary hits for G.

How precisely the discretization identifies the welfare cells. The government's objective is close to flat in its choice of purchases, and the flatness rather than the residual tolerances is what limits the precision of the reported CEVs. Because the flow payoff depends on inherited G_t and not on the choice G', the entire variation of the objective across candidates is the discounted continuation value, which the converged checkpoint stores. Rebuilding that profile exactly as the solver does, at every on-path state on the headline grid, gives a median advantage of the chosen node over its runner-up worth 0.10% of consumption, with about half of cells below a tenth of a percent. The set of candidates within one percent of consumption of the optimum spans about 61% of mean purchases. Refining the discretization makes this worse rather than better: re-allocating the same forty-eight nodes so that the dense block covers the ergodic range, which cuts the local spacing on that range by a factor of more than four, lowers the median advantage from 0.102% to 0.019% and changes the sign of three of the four rule cells. A finer ladder multiplies near-ties instead of resolving them. Section 6.4 shows that this flatness is a property of the equilibrium rather than of the numerical method: the government's objective is flat in G' for the same reason that a one-period deviation from the equilibrium level is worth a thirtieth of a sustained one.

A second, separate source of imprecision is the gap between the stored purchases policy used to simulate the benchmark and the exact discrete best response. On the $A_{\text{ref}} = 0$ companion that gap averages 5.7% of mean purchases, or 1.04 grid steps, and 11.6% of cells differ by more than a full step. Priced with the transform of (29), moving that benchmark to the exact best response is worth +1.78% of consumption, which exceeds the reported rule cells. Because every CEV in Table 6 is computed against the same benchmark simulation, this gap cancels exactly from any comparison between two counterfactuals, and in particular from the state-conditional contrast on which the paper's cyclical conclusions rest. It does not cancel from the level of a cell measured against the benchmark, which is a further reason to read those levels as bounded.

Euler and Bellman residuals are below the convergence threshold. The G grid edges are numerical. The baseline government objective contains no separate tax-deviation term; Section 6.3 reports the robustness specification that adds (26). Productivity uses $n_z = 3$ Rouwenhorst nodes

(Rouwenhorst, 1995; Kopecky and Suen, 2010; Tauchen, 1986); the world-rate and preference chains are those of Section 4. Table 9 records the resulting counts.

Table 9: Computational grids and state counts

Object	Construction	Count
Base exogenous block	(z, r^*, θ)	$3 \times 2 \times 2 = 12$
News block	one-period-ahead (z_{t+1}, r^*_{t+1})	$3 \times 2 = 6$
Full exogenous state	base $\times$ news	$12 \times 6 = 72$
Headline (welfare) (A, G) grid	$n_a \times n_g$	60×48
Re-allocated (A, G) grid (Appendix I)	$n_a \times n_g$	60×48
Tabular policy cells (headline)	full exog $\times\, n_g$	$72 \times 48 = 3{,}456$
Household coefficient groups	base exog only (news pooled)	12
Government coefficient groups	base exog only (news pooled)	12

Notes: Tabular policies are indexed by the 72-node exogenous state at every purchases node (Algorithm 1). The projections in Table 10 pool the news nodes into the 12 base groups. Moments and welfare use the 60×48 tabular pipeline; the re-allocated grid of Appendix I moves only where those nodes sit.

D Numerical algorithm

We solve the Markov-perfect equilibrium by fixed-point iteration on the pair (H, Ψ), following the structure of Bachmann and Bai (2013b) and adapted to a small open economy in which capital is not a state. The objects (H, Ψ) are log-quadratic projections, grouped by the current exogenous block (z, r^*, θ) and pooling the news nodes:

$$\log A' = \alpha_0 + \alpha_1 \log A + \alpha_2 (\log A)^2 + \alpha_3 \log G + \alpha_4 (\log G)^2, \tag{30}$$

$$\log G' = \beta_0 + \beta_1 \log G + \beta_2 (\log G)^2 + \beta_3 \log A + \beta_4 (\log A)^2. \tag{31}$$

The tabular policies that EGM and the discrete government problem produce are the objects the simulation interpolates. The projections (30)–(31) are the maps on which the outer loop iterates.

Algorithm 1 (Fixed-point iteration on (H, Ψ))**. Step 0.** Fix the functional forms (30)–(31). Choose an initial coefficient vector (α^0, β^0) and the implied pair (H^0, Ψ^0). Set the coefficient tolerance $\varepsilon = 10^{-3}$. Build the interpolation grids: $n_A = 60$ asset nodes, $n_G = 48$ purchases nodes (center-dense around G_{ref}), and the 72-node exogenous block (z, r^*, θ) crossed with one-period-ahead news $(z', r^{*\prime})$. The candidate set for G' is the purchases grid, truncated by the ceiling $G' \le \bar{g}\, Y'$.

Step 1. In outer iteration $n \ge 0$, impose Ψ^n in the household problem and solve the parametric dynamic program by the endogenous-grid method on the tax-collapsed budget. This produces the continuation value V^n and the off-path best response $h^n(a, \mathbf{s}, G')$ on the full candidate grid for G'. The on-path slice is $H^n(\mathbf{s}) = h^n(A, \mathbf{s}, \Psi^n(\mathbf{s}))$. Capital is not interpolated: K is recovered from (z, r^*) by no-arbitrage.

Step 2. Without imposing Ψ^n, and taking h^n and V^n as given, compute the government's discrete best response at every state. For each candidate G'_i on the purchases grid, evaluate household value under the implementing tax $\tau_\ell(\mathbf{s}, G'_i)$ and the ceiling. The recorded choice is

$$\Psi^{n,\text{tab}}(\mathbf{s}) \;=\; \arg\max_{G'_i \in \Gamma(\mathbf{s})} J\big(A, \mathbf{s}, G'_i; h^n, V^n\big).$$

Near-ties are broken by a hysteresis bonus of 2×10^{-3} on the incumbent node and, when values remain within 10^{-10}, by the smaller G'. A one-dimensional parabolic refinement between adjacent nodes is accepted only when it raises J.

Step 3. Simulate $N_H = 1$ household for $T = 20{,}000$ periods and discard the first 1,000. Use the same exogenous-shock series in every outer iteration. Along the path, record the on-path pair (A_{t+1}, G'_t) and, at each date, the off-path asset responses $\{A_{t+1,i}\}_{i=1}^{n_G}$ obtained by evaluating h^n at every candidate G'_i. The collected sample therefore has size of order $T(1 + n_G)$, as in Bachmann and Bai (2013b).

Step 4. For each of the 12 current-state groups (z, r^*, θ), estimate $(\hat{\alpha}^n, \hat{\beta}^n)$ by ridge-regularized least squares. The household coefficients $\hat{\alpha}^n$ are updated on the on-path and off-path samples; the government coefficients $\hat{\beta}^n$ are updated on the on-path sample only. The ridge penalties are 0.10

on H and 0.01 on Ψ. A group is left at its previous coefficients if the design has rank below 3 or fewer than 300 on-path observations.

Step 5. If $|\alpha^n - \hat{\alpha}^n| < \varepsilon$ and $|\beta^n - \hat{\beta}^n| < \varepsilon$, go to Step 6. Otherwise set

$$\alpha^{n+1} = \lambda_H \, \hat{\alpha}^n + (1 - \lambda_H) \, \alpha^n, \qquad \beta^{n+1} = \lambda_\Psi \, \hat{\beta}^n + (1 - \lambda_\Psi) \, \beta^n,$$

with $\lambda_H, \lambda_\Psi \in (0, 1]$ chosen by the adaptive mixer, apply Anderson acceleration with memory 5, and return to Step 1. The first iteration uses $\lambda_H = \lambda_\Psi = 0.01$. Alongside the raw coefficient change we monitor an exponential moving average with smoothing 0.20; at acceptance on the headline solve that average is one to two orders of magnitude below ε.

Step 6. Re-fit (30)–(31) on the converged tabular policies and record the R^2. The projection is a summary, not an acceptance test. Moments, impulse responses, and welfare interpolate the tabular policies. A low R^2—0.34 for H and 0.25 for Ψ on the headline grid—is expected when a step-like purchases rule is projected onto a log-quadratic, and is why the simulation does not use (30)–(31).

The household block in Step 1 is a contraction at a fixed purchases rule, which is why the welfare tables re-solve it to 10^{-6} after the outer loop has stopped. The government block in Step 2 is not. Appendix C records the flatness of J in G', and Appendix E the selection rule used on that landscape.

Two implementation details differ from the closed-economy algorithm in Bachmann and Bai (2013b). First, there is no aggregate-capital law of motion to iterate: K_t is pinned by $F_K(K_t, 1) = r_t^* + \delta$. Second, G' is chosen on a finite grid rather than by a continuous line search. The ceiling $G' \leq \bar{g} \, Y'$ is a hard restriction on that grid; the fitted rule (31) smooths over the kink (Section 5).

E Equilibrium selection

We do not claim a uniqueness theorem for the infinite game. The discrete best response leaves a noise floor in the coefficient fit, so more than one numerical fixed point is possible. All numbers in the paper use the checkpoint obtained from Algorithm 1 with a warm-started purchases policy and the household then re-solved to 10^{-6}. The stored purchases policy is close to a one-step exact best response but not identical to it. On the $A_{\text{ref}} = 0$ companion economy the mean absolute gap is 0.026 (5.7% of mean purchases; maximum 0.52), and 11.6% of cells differ by more than one grid step; an earlier selected 60×48 checkpoint has a mean gap of 0.039. A 21×21 diagnostic produces the same tabular purchases policy from cold and warm starts. We have not established that identity on the headline 60×48 grid, and we do not claim local attraction.

F Policy-function coefficients

Fitted policy rules. Table 10 summarizes the log-quadratic policy coefficients grouped over the 12 exogenous combinations of (z, r^*, θ). On the household side, asset persistence is positive in every

group, with mean $\alpha_1 = 1.426$ and substantial cross-state dispersion (s.d. 0.401, range $[0.855, 2.035]$); curvature is negative in every group (α_2 mean -0.145), so the fitted rule is concave in $\log A_t$; and the response to G_t is negative in every group, with mean $\alpha_3 = -2.051$ and s.d. 0.919. Both coefficients are ordered by the state in the same direction: they are largest in absolute value in low-productivity, high-world-rate groups. On the government side, essentially all cross-state variation is in the state-dependent intercept β_0 (mean 0.055, s.d. 0.133, range $[-0.132, 0.346]$), while own persistence and asset sensitivity are smaller by two to three orders of magnitude ($\beta_1 = -0.0037$ and $\beta_3 = 0.0016$ on average). The reason is visible in the tabular rule itself: the equilibrium purchases policy takes only five distinct values over the entire state space, $\{0.229, 0.371, 0.447, 0.522, 0.598\}$, and in the four high-productivity groups it is exactly constant in (A_t, G_t). The fiscal authority adjusts purchases through state-contingent level shifts, not through smoothing of past G or feedback from household wealth. This is a quantitatively sharp version of the qualitative pattern emphasized in Bachmann and Bai (2013b,a).

Table 10: Policy function coefficients by exogenous state group

	Mean	Std. dev.	Min	Max
Panel A: Household asset policy				
Persistence α_1	1.426	0.401	0.855	2.035
Curvature α_2	−0.145	0.049	−0.207	−0.058
Gov't response α_3	−2.051	0.919	−3.494	−0.733
Gov't curvature α_4	−0.149	0.067	−0.253	−0.053
Panel B: Government spending policy				
Intercept β_0	0.05470	0.13328	−0.13185	0.34567
Own persistence β_1	−0.00374	0.00472	−0.01493	0.00000
Own curvature β_2	−0.00027	0.00035	−0.00110	0.00000
Asset response β_3	0.00158	0.00197	0.00000	0.00624
Asset curvature β_4	−0.00150	0.00163	−0.00523	0.00000

Mean elasticities: $\partial \log A'/\partial \log G = -2.051$; $\partial \log G'/\partial \log A \approx 0.0016$.
Notes: The forms are (30)–(31), centered at $a_{\text{norm}} = 0.070$ and $g_{\text{norm}} = -0.860$, and fit on the household re-solved to 10^{-6} under the equilibrium purchases policy (Algorithm 1, Step 6). Statistics are the mean, population standard deviation, minimum, and maximum across the 12 exogenous groups.
In the four high-productivity groups ($z = 1.102$) the purchases rule is constant in (A_t, G_t), so β_1–β_4 are zero by construction. The projection skips those regressions and recovers $\beta_0 = \log G' - g_{\text{norm}}$.
The projection is a summary: R^2 in logs is 0.34 for assets and 0.25 for purchases (mean symmetric absolute percentage error in levels 87% and 7%). The simulation interpolates the table.
The headline solve is accepted at iteration 87 of a 350-iteration cap. At acceptance the household update is $0.843\,\varepsilon$ and the government update $0.053\,\varepsilon$; the EMA counterparts are $0.073\,\varepsilon$ and $0.005\,\varepsilon$. The household residual on a 9×9 test grid is 2.4×10^{-3} in levels and 2.2×10^{-2} in logs.

G Business cycle statistics on the levels basis

Table 11: Business cycle statistics (baseline with news shocks; headline 60×48 grid, ten seeds, $T = 10{,}000$)

	Mean	Std. (level)	CV	Persistence ρ	Corr. with Y
Output Y	1.6099	0.1718	0.1067	0.8429	1.0000
Consumption C	0.5812	0.1082	0.1862	0.7804	0.6673
Government G	0.4998	0.0741	0.1482	0.6757	0.2915
Labor tax τ_ℓ	0.4785	0.0721	0.1507	0.5427	−0.3927
Macro share of $\mathrm{Var}(G)$					
$R^2\big(\log G_t$ on $\log Y_t\big)$					0.117

Notes: Unified tabular simulation pipeline on the headline 60×48 grid: $T = 10{,}000$ with a 500-period burn-in, cross-seed means over ten seeds (same pipeline and grid as the welfare comparison in Section 6.3, which uses twenty seeds). "Std. (level)" equals Mean×CV and matches directly computed standard deviations (e.g., $\sigma(G) = 0.0741$). Persistence ρ and the correlation with Y are computed on the simulated series in *levels*; the model series are stationary around a constant mean, so no detrending is applied in this table. Both differ from the data convention in Table 1, where moments are computed on HP($\lambda = 100$) cyclical components; the model-versus-data comparison in Table 5 is therefore carried out entirely on the HP basis. The macro share is $R^2 = \mathrm{corr}(\log G, \log Y)^2 \approx 0.117$ on this levels basis. It should not be read as the share of the *cyclical* variance of purchases that output explains, which is far smaller: regressing the HP cyclical component of $\log G$ on that of $\log Y$ gives $R^2 = 0.007$. Almost all of the comovement that the levels figure records is low-frequency, and almost none of the year-to-year cyclical variation in purchases is accounted for by contemporaneous output.

H Impulse responses

The results shown so far do not account for timing. To address this, we calculate generalized impulse responses as described by Koop et al. (1996), which are appropriate here because the equilibrium is nonlinear. For each of several hundred simulations, we let the economy reach its ergodic distribution and then split it into two scenarios that share all future shocks. In one scenario, a single Markov chain is shifted to a specified node at the impact date and then evolves from there. By averaging the differences between these paired scenarios, we obtain the response to a one-node transition from a common starting point within each pair. Figure 4 shows the fiscal block.

We do not report output, capital, or investment because they provide no additional information about the model's endogenous propagation. With internationally mobile capital and no adjustment costs, $K_t/L = \left(\alpha z_t/(r_t^* + \delta)\right)^{1/(1-\alpha)}$, so output is a static function of (z_t, r_t^*). Its response therefore reproduces the shock path and is identically zero in the preference column. The model's endogenous propagation operates through the fiscal block and the household's asset position.

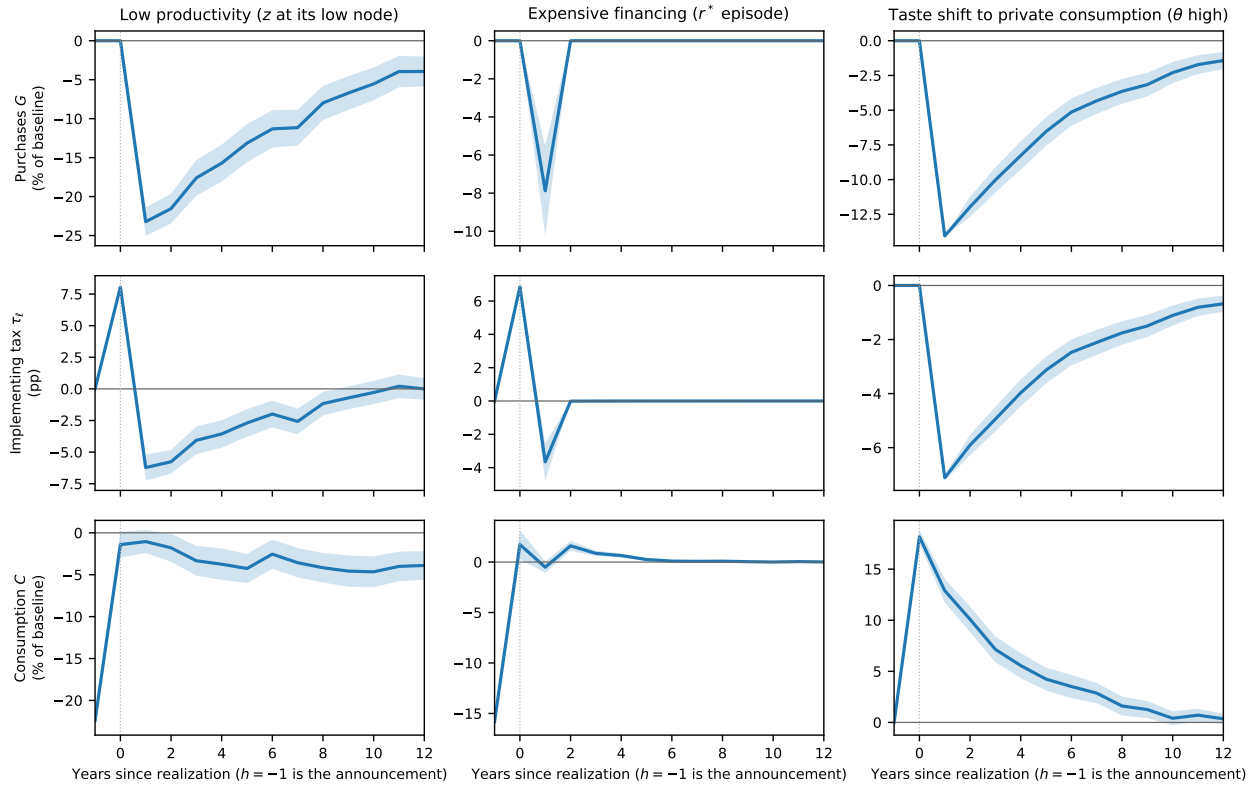


Figure 4: Generalized impulse responses of purchases, the implementing tax, and private consumption.

Notes: Generalized impulse responses (Koop et al., 1996) on the baseline 60×48 solution. Each replication draws a pre-shock history long enough to reach the ergodic distribution and then splits into a baseline and a shocked continuation that share all subsequent innovations; the shocked arm is moved to the named node at $h = 0$ and mean-reverts along the chain thereafter. Responses are averaged over 207, 362, and 202 paired draws for the three columns respectively, and shaded regions are ± 2 standard errors of the mean response. G and C are in percent of the baseline path; τ_ℓ is in percentage points. Because (z_{t+1}, r^*_{t+1}) are known one period in advance, $h = -1$ is the announcement date and $h = 0$ the realization date; there is no news on θ. Output, capital, and investment are static functions of (z_t, r_t^*) and are omitted; their responses reproduce the closed-form shock and are exactly zero under θ.

The timing rule from Section 4 sets G_t at $t-1$, so purchases cannot change in the same period as a shock. When the productivity shock occurs, output drops by 13.5% at $h = 0$, purchases remain unchanged, and the tax increases by 8.1 percentage points to finance the predetermined level of G from a smaller tax base. In the following period, purchases fall by 23.2% and the tax drops 6.2 points below baseline. A similar pattern appears after a world-rate shock but lasts only two periods because expensive-financing episodes are i.i.d. rather than persistent: output is 10.2% below baseline for one period, the tax initially rises by 6.9 points and then falls by 3.7 points at $h = 1$ as purchases decrease by 7.9%, and all series return to baseline by $h = 2$. This illustrates the financing assignment from Section 6.3 over time. Spending adjusts after a one-period delay, and the tax covers what cannot be postponed.

The preference column shows the composition margin. When preferences shift toward private consumption, consumption rises by 18.2% on impact and purchases decrease by 14.1% beginning at $h = 1$. The tax falls by 7.1 points because the smaller level of G requires less revenue. Output remains unchanged because θ affects neither the production function nor the capital rule. The response reallocates a given resource path between C and G, so the preference state changes the level of G/Y and contributes to $\mathrm{Var}(G)$ but has little effect on $\mathrm{corr}(G, Y)$.

The announcement date shows how the news structure creates anticipation. Because (z_{t+1}, r^*_{t+1}) are known at time t, households adjust their saving a full period before output changes. Consumption falls by 22.4% before the productivity shock and by 15.8% before the world-rate episode, while assets rise by 0.120 and 0.092 respectively and return to baseline by $h = 1$. This advance saving leaves consumption only 1.4% below baseline at $h = 0$, even though output falls by 13.5%: households pre-fund the period in which purchases cannot adjust and the tax is highest. These large anticipatory movements result from the polar news assumption rather than from the fiscal rule. The model reveals the entire innovation one period early, so the consumption response is concentrated at the announcement rather than spread over the episode. With a weaker information structure, the response at $h = -1$ would be smaller and the impact response would be larger, but the assignment between spending and taxes would remain the same. As expected, the preference column, which has no news, is flat at $h = -1$.

Three patterns from these results matter for the welfare analysis that follows. First, the procyclicality of G in this model is supply-driven. The variance decomposition attributes most of the systematic component of $\mathrm{Var}(G)$ to productivity and world-rate states. Preference shocks mainly alter the allocation between public and private consumption and have a smaller effect on $\mathrm{corr}(G, Y)$. Second, the policy-coefficient tables show that the rule adjusts mainly through state-dependent levels rather than through high own-persistence in G. Purchases rise in good states because fundamentals improve, not because the government extrapolates past G or reacts strongly to household balance sheets. Third, the impulse responses show that timing matters as much as magnitude. Purchases cannot change when a shock occurs, so the tax absorbs the immediate impact and spending adjusts one year later. The time-consistent rule shifts G with the state. Whether it should also be smoothed over the cycle is the question addressed by the counterfactual.

I Robustness

I.1 Identification and limitations

Ergodic levels and consumption-equivalent magnitudes. Strict impatience implies an ergodic consumption share far below the data ($C/Y \approx 0.36$ against about 0.73) and an extreme external position (NFA of about -286% of output). Under log utility, marginal utility is proportional to $1/C$, so consumption-equivalent *levels*—including permanent proportional restraint and state-conditional cells—are evaluated at roughly twice the marginal utility the data would imply. We therefore report CEV magnitudes as calibrated illustrations under this economy, not as data-consistent welfare percentages. The gap is about impatience rather than asset-grid placement: moving the reference point to $A_{\text{ref}} = 0$ leaves the ergodic C/Y and NFA discrepancies essentially unchanged. A debt-elastic premium on household assets is the standard device for pinning down the external position, but it is not a free correction here. Under strict impatience the household's Euler equation comes to rest where the premium offsets the impatience wedge, so the ergodic position is determined jointly by the premium schedule and the discount factor rather than by the premium's target alone. Matching the position in the data therefore means recalibrating the household block, which alters the interest rate faced in every state and hence the financing elasticities behind Lemma 1. We leave that extension to future work, and read the ergodic level discrepancy as a structural feature of the present calibration rather than as a modeling detail that a small friction would absorb. What the paper compares, however, is CEV magnitudes computed against a common benchmark path at a common marginal utility, which is why the *ranking* of level against cyclical effects, the reassignment result, and the opposite state-conditional incidence of the two budget sides are on firmer ground than any single CEV level.

Expenditure ceiling and cyclical purchases. The baseline ceiling $\bar{g} = 0.35$ binds at the discrete best response in a large share of visited states, and realized G_t/Y_t exceeds it in about 27% of periods. Raising the ceiling to $\bar{g} = 0.50$ and re-solving leaves the cyclical sign of purchases intact: across the pair of economies that differ only in the ceiling, HP $\text{corr}(\tilde{G}, \tilde{Y})$ moves from $+0.08$ to $+0.06$ and relative volatility remains above one, falling from 1.96 to 1.70. Both economies in this pair are solved at $A_{\text{ref}} = 0$, whose $\bar{g} = 0.35$ relative volatility of 1.96 is that of the headline economy to within rounding. Equilibrium procyclicality is therefore not an artifact of the tight ceiling alone. The raised-ceiling economy does shift the ergodic level of G/Y sharply upward, from 0.32 to 0.48, so the comparison speaks to the cyclical sign without claiming level invariance. Imposing the ceiling symmetrically on exogenous-rule paths can move the spending-side cyclical CEV by several points, which is why the ergodic cyclical cells are quoted as a band.

Permanent restraint across normalizations and grids. Permanent proportional restraint peaks at $+13.9\%$ of consumption when $A_{\text{ref}} = K_{\text{ref}} + \bar{B}$ (mean G/Y from 0.312 to 0.172), with a one-period deviation of about -0.47% by the government's own objective. The peak is $+15.1\%$

under $A_{\text{ref}} = 0$ and about +20.1% when the same number of purchases nodes is re-allocated over the ergodic range of G. In each case the one-period arm remains of order one half of one percent, so the level-versus-cycle ranking and the large gap between sustained and one-period restraint survive the discretization change that flips three of four *cyclical* rule signs in Table 12. The point estimate moves with the ergodic purchases level; the qualitative identification of the level gap does not.

Other numerical gaps. The mean absolute gap between the stored purchases policy and a one-step exact best response is about 5.7% of mean G on the $A_{\text{ref}} = 0$ companion (about 1.8% of consumption if priced as a CEV against that path). Lemma 1 is reported at $\bar{r}^* = 4\%$; at the modal visited node $r^{*,L} = 2.48\%$ the world-rate elasticity is about -10 rather than -7.1. The -12.6% spending-side cell in the (low z, high r^*) bin has a cross-seed standard deviation of about 0.6 percentage points.

I.2 Grid and normalization robustness

The headline welfare table uses the 60×48 grid (Table 6). Using the same protocol—household-resolve, ergodic-mean CEV, $T = 10{,}000$, burn-in 500, and the same 20 seeds for each row—we recompute all four rules under two variations of that solve: the alternative asset normalization, and a re-allocation of the same number of purchases nodes over the ergodic range. Table 12 reports all three, each calculated using the unified simulation pipeline described in Section 5.1.

Reading the table. The first two rows answer the normalization question and the third answers the discretization question, and they answer them differently.

Moving the reference point changes almost nothing. Dropping the household's claims on capital and public debt ($A_{\text{ref}} = 0$) moves the reference labor tax from 0.219 to 0.243 and the reference external position from zero to -268% of output, yet the counter-cyclical purchases cell is identical to two decimals, the two constant-level cells move by 0.20 and 0.04 points, and the largest movement, in the counter-cyclical tax cell, is 0.85 points against cross-seed standard deviations of about 0.2. The ergodic moments are equally insensitive: mean G/Y moves from 0.313 to 0.319 and the HP

Table 12: Financing-assignment CEVs under three discretizations of the same economy

Grid	$\text{CEV}_{A.i}$ (%)	$\text{CEV}_{A.ii}$ (%)	$\text{CEV}_{B.i}$ (%)	$\text{CEV}_{B.ii}$ (%)	Seeds
60×48 headline, $A_{\text{ref}} = K_{\text{ref}} + \bar{B}$	-1.20	-3.11	-0.47	-1.59	20
60×48, $A_{\text{ref}} = 0$	-1.20	-3.31	-1.32	-1.63	20
60×48, nodes re-allocated	0.48	-0.79	1.65	1.62	20

Notes: All rows use the common protocol (household-resolve, ergodic-mean CEV, $T = 10{,}000$, burn-in 500, one-pass mean G/Y scaling at seed 4242), the unified simulation pipeline (Section 5.1), and the same 20 seeds. Cross-seed standard deviations range from 0.21 to 0.58 percentage points. A.i is counter-cyclical G (τ_ℓ residual); A.ii is constant G; B.i is the output-indexed tax rule (G residual); B.ii is constant τ_ℓ. The first two rows differ only in the reference asset normalization of Table 8. The third row solves the same economy as the first with the same number of purchases nodes, re-allocated so that the dense block covers the ergodic range of G rather than the neighbourhood of G_{ref}; this cuts the local spacing by a factor of more than four. Ergodic cyclical cells are reported as a band (Section I.1); Source: replication package.

correlation of purchases with output stays in $\{+0.08, +0.09\}$. The welfare comparisons in this paper are therefore not artifacts of the asset normalization, which is the most economically consequential of the calibration choices behind the reference point.

Re-allocating the grid nodes changes a great deal. Holding the economy, the protocol, the seeds and the node count fixed, and moving only where the nodes sit, flips the sign of three of the four cells and moves the counter-cyclical purchases cell from -1.20 to $+0.48$. The re-gridded equilibrium is not defective: its policy is interior, its household re-solves converge, and its cyclical moments are close to the headline economy's, with $\text{corr}(\tilde{G}, \tilde{Y}) = +0.08$ against $+0.09$. What moves is the ergodic level, with mean G/Y at 0.339 against 0.313, and the welfare cells follow the level. This is the flatness documented in Appendix C seen in its consequences, and it is why we report the ergodic cells as bounded and rest the paper's conclusions on the state-conditional contrast, which is an order of magnitude larger, and on the level result of Section 6.4, which is larger still.

The variations reported here are deliberately not a grid ladder. A second grid built the same way as the first, with the dense block placed around G_{ref}, would refine the spacing without addressing the question the third row answers: a finer ladder placed the same way multiplies near-ties rather than resolving them, which is the pattern Appendix C documents directly. Re-allocating the nodes is the informative variation, and it is the one we report.

What the table establishes is therefore narrower than convergence of the cells. It establishes that the level result and the state-conditional contrast are invariant to both variations, and that the ergodic magnitudes of the four cyclical cells are not identified to better than the band we quote for them.

Robustness to including the adjustment cost in Panel B. The Panel B closure $G_t = \tau_{\ell,t} w_t L + \tau_a r_t^* A_t - r_t^* \bar{B}$ omits the strategic adjustment cost on G when the τ_ℓ rule is exogenous. We test this convention using an alternative Panel B in which $AC_t = (\Omega/2)(G_t - G_{t-1})^2$ is computed along the simulated rule path and *deducted* from the household's effective public-good consumption; that is, utility uses $\max\{G_t - AC_t, \varepsilon\}$ rather than G_t. This treats AC_t as a passive resource cost paid mechanically by the government even though the rule is exogenous.

The implied adjustment cost is extremely small under both Panel B rules. Average AC_t/G_t is about 6×10^{-6} for the countercyclical-τ_ℓ rule and 5×10^{-6} for the constant-τ_ℓ rule.

Consequently, the AC-adjusted CEV is indistinguishable from the no-AC CEV: the difference is less than 10^{-3} percentage points in either Panel B cell. The treatment of the adjustment cost across the two panels therefore does not drive the welfare ranking, and the state-conditional contrast between the spending-side and financing-side assignments is unchanged whether or not this passive adjustment cost is charged under the exogenous τ_ℓ rule.

Diagnostics. For the household object used in the main welfare calculations, the re-solve under the baseline purchases policy converges after 23 iterations. The final value-function change is 5.96×10^{-8} and the largest change in the equilibrium asset policy is 8.49×10^{-7}, both below the 10^{-6} stopping tolerance. The simulation does not reach any G-grid boundary (see Section 5.1),

and the fraction of periods with $A_t \leq 0$ in the baseline simulation is zero. This means the log basis used in the projection is evaluated only on the positive-A region where the ergodic distribution is located. Along with the moment match in Table 5 and the cycle decomposition in Section 6.2, these checks show that the main welfare comparison is not affected by nonconvergence of the household re-solve. However, they do not confirm the accuracy of the fitted purchases projection near the expenditure ceiling described in Section 5. Since the benchmark is simulated from the converged equilibrium policy and the counterfactuals use outside rules, any projection slack appears only in the benchmark part of each comparison, not in the counterfactual part. Three points limit how much this slack could matter. First, the slack is both localized and moderate: it only appears in high-spending-demand states where the fitted rule smooths over the ceiling kink, with realized G/Y reaching up to 0.48 while the ceiling is 0.35. Second, every CEV in Table 6 is measured against the same benchmark path, so the slack cancels out in any comparison between counterfactuals. This means the contrast between the counter-cyclical and constant purchases rules, and comparisons among the four rules, do not depend on it. The slack can only affect the level of each CEV compared to the benchmark, and might change the sign of a cell if its value is small. Third, the two findings the paper rests on are of a kind the slack cannot generate. The state-conditional contrast of Section 6.3 is a comparison between two counterfactuals and so falls under the previous point, and the level result of Section 6.4 is measured between two scalings of the equilibrium's own rule rather than against the fitted benchmark at all. For these reasons, we see the slack as a remaining uncertainty about CEV levels, especially for the less certain Panel B values, but not as something that changes the main findings of the paper.