WP-EMS

Working Papers Series in Economics, Mathematics and Statistics

1506
UNIVERSITÀ
DEGLI STUDI
DI URBINO
CARLO BO

WP-EMS # 2026/02

The Political Economy of Tax Expenditures:

Fragmentation versus Volume

Mauro Marè
Luiss University
Francesco Porcelli
Sapienza University of Rome
Francesco Vidoli
University of Urbino

September 14, 2026

ISSN 1974-4110 (on)
ISSN 1594-7645 (pr)

www.econ.uniurb.it

The Political Economy of Tax Expenditures

Fragmentation versus Volume

Mauro Marè[a], Francesco Porcelli[b,*], Francesco Vidoli[c]

[a] *Luiss University*
[b] *Sapienza University of Rome*
[c] *University of Urbino Carlo Bo*

Abstract

Tax expenditures are delivered through many legally distinct provisions rather than through a few salient programs. We argue that the number of provisions, not the aggregate revenue foregone, is the margin on which their political economy operates. We develop a three-agent model in which a rent-seeking firm lobbies for tax expenditures, a vote-maximizing politician chooses visible public goods, the volume of tax expenditures and the number of provisions through which they are delivered, and a Treasury official with an institutional preference for budget transparency can expose their hidden cost at the risk of political retaliation. Fiscal illusion is micro-founded as an increasing, concave function of the number of provisions, so that fragmentation is the channel through which opacity is produced. In the subgame-perfect equilibrium, a stronger Treasury enforces more, and stronger enforcement reduces both the volume and the fragmentation of tax expenditures while leaving visible public goods essentially unchanged. The number of provisions has no source of variation independent of the opacity mechanism, whereas revenue foregone also reflects legitimate redistributive demand: the count of provisions is therefore a structurally cleaner signal of political capture than the aggregate fiscal magnitude, and institutional quality should discipline fragmentation more robustly than volume. The model yields three testable hypotheses for cross-country data in which provisions are counted separately from revenue foregone.

Keywords: Tax Expenditures, Fragmentation, Fiscal Illusion, Lobbying, Budget Institutions
JEL: D72, D78, H20, H30, H61

1. Introduction

Governments deliver a large and growing share of fiscal policy through the tax system. Tax expenditures, the revenue foregone through deductions, exemptions, credits and reduced rates that deviate from a benchmark tax structure, average more than three percent of

*Corresponding author. Department of Economics and Law, Sapienza University of Rome. E-mail: francesco.porcelli@uniroma1.it. An earlier version of this work was presented at the 82nd Annual Congress of the International Institute of Public Finance (ISEG, Lisbon, August 2026); we thank Andrea Tulli and the session participants for helpful comments.

GDP among the jurisdictions that report them, and in several countries they are delivered through hundreds of legally distinct provisions (Redonda et al., 2023). The public-finance literature has long recognized that tax expenditures are spending programs delivered off-budget (Surrey, 1973; Kleinbard, 2010; **?**), that they escape the scrutiny of the appropriation process (Burman and Phaup, 2012), and that their political appeal lies in their low salience (Tyson, 2014). What the literature has not done is to treat the two dimensions along which tax expenditures can be measured, their aggregate cost (*volume*) and their splitting into distinct provisions (*fragmentation*), as separate choices, and to ask which of the two responds to the institutions meant to discipline hidden redistribution.

The model in words. Three players interact. A *firm* wants tax breaks that lower its costs and spends resources on lobbying to obtain them; the more effective its lobbying and the cheaper political influence, the larger the tax break it demands. A *politician*, who cares about votes, decides how much to spend on visible public goods, how large a tax expenditure to grant, and, crucially, through how many separate provisions to deliver it. A *Treasury* official, who cares about budget transparency, can quantify and publicize the hidden cost of tax expenditures, at the risk of being removed by a politician who resents the exposure.

The engine of the model is a single idea about attention: voters can only follow so many things. A large, single tax break is noticed and its cost is debated like any spending item; the same amount split into dozens of narrow provisions, each with its own beneficiaries and its own technical justification, is not. Opacity is therefore not a fixed feature of the tax system but something the politician produces by fragmenting tax expenditures, and it is worth producing because it lowers the electoral cost of rewarding the firm. Fragmentation is not free, however: each provision must be drafted, defended and coordinated, and each new loophole draws some scrutiny. The politician fragments up to the point where the electoral cost saved by the next provision equals its legislative cost. The Treasury's enforcement enters here: by exposing part of the hidden cost, it makes each provision less effective at hiding, and so lowers both the tax expenditure the politician is willing to grant and the number of provisions used to deliver it. Visible public goods, which are fully salient, are not the margin of adjustment.

Two consequences follow, and they are the paper's message. First, fragmentation exists for one reason only, to hide tax expenditures, and that reason is exactly what the Treasury's enforcement attacks: the number of provisions falls with the institutional strength of the Treasury. Second, the aggregate cost of tax expenditures is generated, in addition to this opacity-driven mechanism, by legitimate redistributive demand, the size of the welfare state, demographic structure and the policy goals delegated to the tax system, which is outside the political-economy mechanism and need not fall with institutional strength. Volume is a noisy signal of capture; fragmentation is a clean one.

Results. We formalize this in a three-stage game solved by backward induction. The firm's demand for tax expenditures rises with lobbying effectiveness and falls with the cost of rent seeking. The politician's optimal fragmentation equates the marginal fiscal illusion generated by an additional provision to its marginal legislative cost, and visible public goods are approximately independent of enforcement. The Treasury's enforcement is strictly increasing in the ratio between its institutional weight on transparency and the political cost of removal. The main proposition establishes that a stronger Treasury enforces more, and that

stronger enforcement reduces both the volume and the fragmentation of tax expenditures while leaving visible public goods essentially unchanged; the corollary states the Volume-Fragmentation Asymmetry: the count of provisions is a structurally cleaner empirical signal of opaque political capture than revenue foregone, so that cross-country associations between institutional quality and fragmentation should be larger and more robust than those with volume. A numerical illustration displays the joint comparative statics of the share of tax expenditures over structural opacity and Treasury strength.

Section 2 relates the paper to the literature and Section 3 presents the stylized facts. Section 4 sets up and solves the model and states the main results. Section 5 derives the testable hypotheses, and Section 6 concludes. Derivations and proofs are in Appendix A; a one-dimensional slice of the numerical illustration is in Appendix B.

2. Related Literature

The paper connects four strands of research.

Tax expenditures. Since Surrey (1973) the public-finance literature has treated tax expenditures as spending programs delivered through the tax code, economically equivalent to appropriations but exempt from their procedural discipline (Kleinbard, 2010; **?**; Burman and Phaup, 2012). Measurement issues, from the choice of the benchmark to the aggregation of revenue foregone, have received considerable attention (Villela et al., 2010; Tax Administration Research Centre, 2014; Burman et al., 2008), and the Global Tax Expenditures Database has recently made cross-country comparison possible (Redonda et al., 2023). Tyson (2014) provides a political-economy account of why legislators prefer tax expenditures to direct spending, based on their lower salience. Our contribution is to model the number of provisions as a choice variable distinct from the aggregate cost, to show that it is the margin on which opacity is produced, and to derive the asymmetric response of the two margins to budget institutions.

Special-interest politics. The lobbying firm of our model belongs to the rent-seeking and pressure-group tradition of Tullock (1967), Krueger (1974) and Becker (1983), and to the common-agency framework of Grossman and Helpman (1994, 2001), in which organized interests set the marginal political value of policy and politicians respond; the textbook treatment is Persson and Tabellini (2000). Relative to that literature, in which the politician chooses how much to concede, our politician also chooses through how many instruments to concede, and faces an institutional counterpart, the Treasury, with an objective of its own. The Treasury is a conservative agent in the sense of Rogoff (1985), whose mandate differs from social welfare, and the contrast between its incentives and the politician's follows the bureaucrat-politician literature of Niskanen (1971) and Alesina and Tabellini (2007); the game structure, in which policy is the equilibrium of a conflict among agents with different objectives, mimics the theory of privatization of Shleifer and Vishny (1994), Boycko et al. (1994) and Boycko et al. (1996).

Fiscal illusion and salience. The idea that voters misperceive the fiscal burden goes back to Puviani (1903) and was revived by Buchanan and Wagner (1980) and surveyed by Oates (1988). The recent literature on salience has documented that tax-based instruments are less

visible than spending and that visibility shapes behavior (Chetty et al., 2009; Finkelstein, 2009), and has modeled limited attention as a scarce resource allocated across items (Bordalo et al., 2013; Gabaix, 2014). We use this mechanism to give fiscal illusion an operational microfoundation: structural opacity is a function of the number of provisions, and it is therefore produced, and disciplined, in equilibrium.

Budget institutions and transparency. A large literature relates fiscal outcomes to the design of the budget process, fiscal rules, transparency and fiscal councils (Alesina and Perotti, 1996; von Hagen and Harden, 1995; Hallerberg et al., 2004; Alt and Lassen, 2006; Wyplosz, 2013; Debrun and Kinda, 2017), and to the informal institutions that support fiscal cooperation, in particular generalized trust and civic capital (La Porta et al., 1997; Tabellini, 2008; Aghion et al., 2010; Guiso et al., 2011). Milesi-Ferretti (2004) shows that rules which target a headline aggregate invite creative accounting. Our Treasury's institutional weight summarizes both the formal and the informal side of this literature, and our asymmetry result is the tax-expenditure counterpart of the creative-accounting concern: the margin on which hidden redistribution escapes a headline rule is fragmentation.

3. Two Margins of Tax Expenditures: Stylized Facts

Tax expenditures have traditionally been measured by a single metric, the aggregate revenue foregone as a share of GDP. Yet they are rarely delivered as a single program: they take the form of many legally distinct provisions, each targeting a narrow base, a sector or an activity, and the count of provisions is a second, independent dimension of the phenomenon. The Global Tax Expenditures Database (GTED), compiled by the Council on Economic Policies and the German Institute of Development and Sustainability (Redonda et al., 2023), records both dimensions for a large set of jurisdictions. Three facts stand out.

Fact 1: fragmentation is large and rising. In the GTED-reporting jurisdictions observed over 2010–2022, the number of distinct provisions ranges from a few dozen to more than one thousand, with a mean of about 190; revenue foregone averages 3.4 percent of GDP. Figure 1 shows the evolution of the two margins by continental group. The aggregate cost is regionally heterogeneous and, outside North America, moves slowly; the average number of provisions displays a clear upward trend in almost all groups. Governments increasingly slice benefits into narrower concessions even where the aggregate cost is stable.

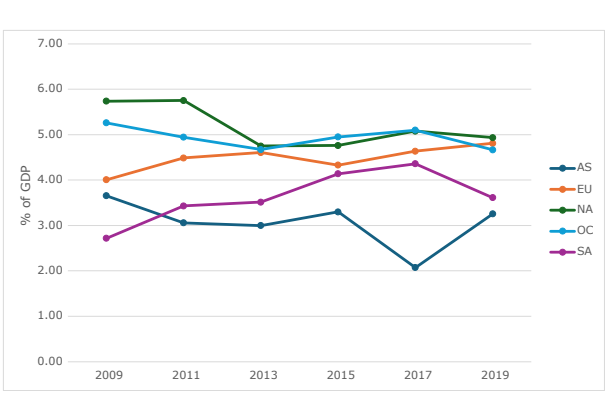

(a) Revenue foregone (% of GDP)

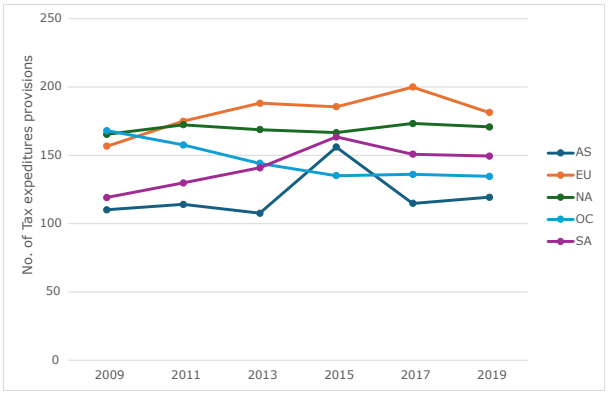

(b) Number of provisions

Figure 1: The two margins of tax expenditures by continental group, GTED-reporting jurisdictions, 2010–2022. Source: authors' elaboration on GTED.

Fact 2: the two margins identify different countries. The jurisdictions that stand out on the volume margin are not those that stand out on the fragmentation margin. In the 2019 cross-section, a few countries (Russia, the Netherlands) report revenue foregone above twelve percent of GDP with a moderate number of provisions, while Italy and France exceed four and five hundred provisions respectively with an aggregate cost that is far less extreme. A given fiscal cost can be the residue of a few broad, salient programs or of hundreds of narrow ones.

Fact 3: the two margins are only weakly correlated. Figure 2 plots the number of provisions against revenue foregone in 2019. The correlation is positive but weak, and a substantial cluster of countries sits in the high-fragmentation, moderate-cost region. Fragmentation is not a mechanical by-product of size.

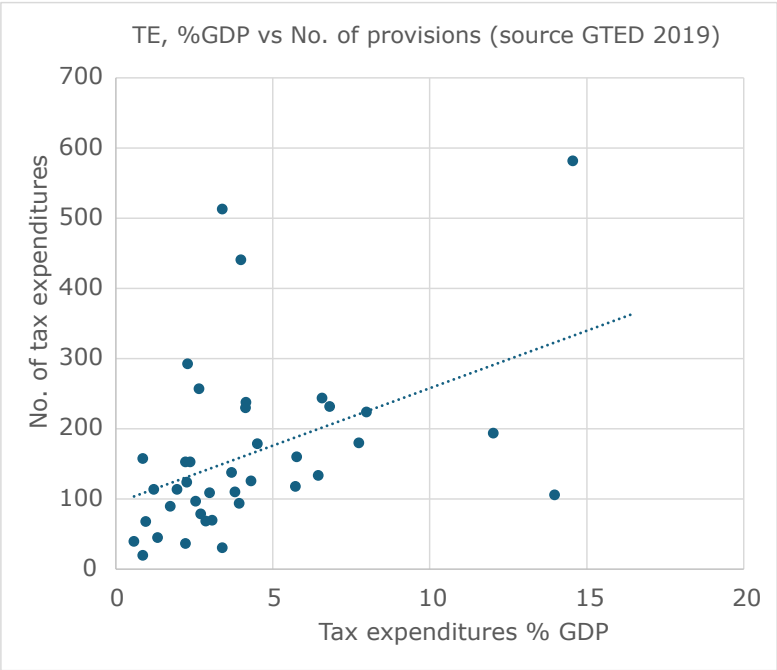

Figure 2: Volume versus fragmentation, GTED-reporting jurisdictions, 2019. Source: authors' elaboration on GTED.

If tax expenditures are attractive because their cost is less salient than spending, what determines how a given cost is split into provisions, and which margin do budget institutions discipline? The model that follows treats fragmentation as a choice and answers both questions.

4. A Three-Agent Political Economy Model of Tax Expenditures

4.1. Model Setup and Theoretical Framework

We develop a political economy model with three strategic agents: a rent-seeking lobbying firm, a vote-maximizing politician, and a budget-conscious Treasury official. The novelty relative to the existing political-economy literature on taxation is twofold. First, the politician is endowed with two strategic margins over tax expenditures: their aggregate volume E and the number N of legally distinct provisions through which E is delivered. Second, fiscal illusion is micro-founded as an increasing function of N, with $\delta = \delta(N)$, so that fragmentation is the operational channel through which structural opacity is generated. This framework extends the traditional public choice literature (Tullock, 1967; Buchanan and Wagner, 1980; Becker, 1983) by focusing on the hidden nature of tax expenditures as a mechanism for political redistribution.

We formalize structural opacity through a function $\delta(N) \in [0, \delta_{\max}]$ with $\delta_{\max} \leq 1$, satisfying $\delta(0) = 0$, $\delta'(N) > 0$, $\delta''(N) < 0$, and $\lim_{N\to\infty} \delta(N) = \delta_{\max}$. The parameter $\delta_{\max}$ reflects the maximum salience differential between tax-based and expenditure-based fiscal interventions documented by Chetty et al. (2009) and Finkelstein (2009); when $\delta_{\max} = 0$ tax expenditures are as transparent as visible public goods and no fiscal illusion is possible,

when $\delta_{\max} = 1$ they are maximally opaque. Following Burman and Phaup (2012), this informational asymmetry creates systematic biases in favour of tax expenditures over programs financed through visible public goods.

The fragmentation micro-foundation of δ. The structural opacity δ should not be interpreted as a fixed environmental constant, but as the reduced-form image of a more granular policy choice. Governments do not implement tax expenditures as a single, salient program: they distribute them across a large number of narrowly targeted legislative provisions. Following the literature on bounded attention and salience (Chetty et al., 2009; Finkelstein, 2009; Bordalo et al., 2013), the salience of any individual tax preference is decreasing in the number of distinct provisions competing for voter attention. A reduced-form specification consistent with the properties of $\delta(\cdot)$ stated above is

$$\delta(N) \;=\; \delta_{\max}\big[1 - 1/(1+\nu N)\big], \qquad \nu > 0, \tag{1}$$

where N denotes the number of legislatively distinct tax expenditure provisions, ν measures the marginal opacity yield of each additional provision, and the concavity of $\delta(\cdot)$ reflects decreasing returns to fragmentation, with each additional narrow loophole adding opacity at a diminishing rate.

This formalization has three implications that anchor the subsequent empirical analysis. First, the aggregate volume of tax expenditures ($E \cdot q^*$) and their fragmentation (N) are *conceptually distinct* margins of the politician's strategy: the same total fiscal cost can be delivered through a small number of broad, salient provisions (low N, low δ) or through hundreds of narrow ones (high N, high δ). Second, N is the empirically observable counterpart of the latent parameter δ through the micro-foundation $\delta = \delta(N)$. Third, the political-economy mechanism developed in this paper operates simultaneously on both margins through the politician's joint choice of E and N formalized in Section 4.3.

By mimicking the approach adopted in the theory of privatization by (Shleifer and Vishny, 1994; Boycko et al., 1994, 1996), we model the political process around tax expenditures as a subgame-perfect Nash equilibrium involving rent-seekers, politicians and the Treasury (the Minister of the Economy and Finance), which is responsible for the public budget and the stability of public finances. Unlike previous models, we treat all three agents as strategic players whose decisions are jointly determined in equilibrium. The Treasury is not merely a passive enforcer of a budget constraint, but an active player whose institutional strength affects the equilibrium allocation between visible public goods and opaque tax expenditures.

4.2. Agent 1: The Lobbying Firm

The lobbying firm operates in a market with inverse demand function $p(q) = A - bq$, where $A > 0$ represents the demand intercept and $b > 0$ the slope parameter. Tax expenditures E reduce the firm's effective marginal cost from c to $(c - \theta E)$, where $\theta > 0$ measures the *lobbying effectiveness*, the ability to convert political favors into cost reductions for the firm. Obtaining these tax breaks is not costless: the firm engages in rent-seeking activities at quadratic cost $R(E) = (\mu/2)E^2$, where $\mu > 0$ captures the intensity of political competition for tax favors (Krueger, 1974). This micro-foundation ensures that tax expenditures are not pure transfers among interest groups, but involve dead-weight losses.

The firm chooses output q and the level of tax expenditures E it demands by solving

$$\max_{q,E} \pi = (A - bq)q - (c - \theta E)q - \tfrac{\mu}{2}E^2. \tag{2}$$

Solving the system of first-order conditions (derived explicitly in Appendix A.1) yields the firm's optimal output and equilibrium demand for tax expenditures:

$$q^* = \frac{\mu(A-c)}{2b\mu - \theta^2}, \qquad E_d^* = \frac{\theta(A-c)}{2b\mu - \theta^2}. \tag{3}$$

The two expressions deliver three economic intuitions. First, the firm's production is increasing in the lobbying effectiveness θ, but the marginal benefit of lobbying is dampened by the rent-seeking cost parameter μ. When μ is large (intense political competition), the firm must spend more resources to obtain the same tax break, reducing the net benefit. Second, equilibrium demand for tax expenditures inherits the same comparative statics: E_d^* rises with θ and falls with μ. Third, we maintain throughout the analysis the stability condition

$$2b\mu > \theta^2, \tag{4}$$

which ensures that rent-seeking costs grow faster than benefits and prevents unbounded lobbying.

4.3. Agent 2: The Politician

The politician chooses simultaneously the level of visible public goods G, the level of tax expenditures E, and the number of distinct legislative provisions N through which E is delivered. Her utility function combines a quadratic specification for (G, E, T) with an explicit political-legislative cost of fragmentation $\psi(N)$:

$$V(G, E, N, T) \;=\; \alpha G - \tfrac{1}{2}k_G G^2 + \beta E - \tfrac{1}{2}k_E E^2 - \gamma T - \tfrac{1}{2}k_T T^2 - \psi(N), \tag{5}$$

where $\alpha, \beta > 0$ are marginal political benefits, $\gamma > 0$ is voter tax aversion, and $k_G, k_E, k_T > 0$ are concavity parameters. The function $\psi(N)$ is strictly increasing and strictly convex ($\psi'(N) > 0$, $\psi''(N) > 0$, $\psi(0) = 0$): it captures the political-legislative cost of proliferating provisions (parliamentary drafting time, attention-budget per provision, residual exposure to press and fiscal-council scrutiny, marginal coordination cost of an increasingly granular tax code). Appendix Appendix A.2 uses the tractable parametrisation $\psi(N) = (\psi_0/2)N^2$ for explicit derivations; the substantive results do not rely on this specific form.

Throughout the derivations we adopt the working assumption that the quadratic component of tax aversion is small relative to the linear component ($k_T \to 0$), consistent with empirical estimates of voter responsiveness to tax burdens in advanced democracies. The politician's first-order conditions (full derivation in Appendix Appendix A.2) deliver

$$\begin{aligned} G^* \;&\approx\; \frac{\alpha - \gamma}{k_G}, \\ E^* \;&\approx\; \frac{\beta - \gamma\,\Omega(N^*, \xi^*)\,q^*}{k_E}, \\ \psi'(N^*) \;&=\; \gamma\,\delta'(N^*)\,(1 - \xi^*)\,E^*\,q^*, \end{aligned} \tag{6}$$

where $\Omega(N, \xi)$ is the *effective visibility* of tax expenditures defined in the next subsection. The third condition pins down equilibrium fragmentation N^*: it equates the marginal political benefit of an additional provision (the extra fiscal illusion generated, weighted by γ and by the equilibrium volume $E^* q^*$) to its marginal political-legislative cost $\psi'(N^*)$. Under the maintained assumptions on $\delta(\cdot)$ and $\psi(\cdot)$, the equation has a unique interior solution $N^* > 0$ whenever the marginal benefit at $N = 0$ exceeds $\psi'(0)$. Equation (6) also shows that G^* is approximately independent of Treasury enforcement under the small-k_T approximation: the substitution between E and G as the Treasury's strength rises is therefore quantitatively driven by the contraction of E^*, not by an expansion of G^*.

4.4. Agent 3: The Treasury

The Treasury official seeks to discipline the politically-exploited share of structural opacity while avoiding removal by the politician. Define the *effective visibility* of tax expenditures as

$$\Omega(N, \xi) \;\equiv\; 1 - \delta(N)(1 - \xi), \qquad \Omega(N, \xi) \in [1 - \delta(N),\; 1]. \tag{7}$$

At $\xi = 0$, the Treasury exerts no enforcement and a fraction $\delta(N)$ of the fiscal cost of tax expenditures is hidden, so $\Omega = 1 - \delta(N)$; at $\xi = 1$, the Treasury fully exposes the hidden portion and $\Omega = 1$. The politician's perceived tax burden is then $T = G + \Omega(N, \xi)\, E\, q^*$, and the *fiscal illusion gap*, the difference between actual revenue requirement and perceived burden, is

$$\Delta_{FI}(N, \xi) \;\equiv\; \left[G + E\, q^*\right] - T \;=\; \delta(N)(1 - \xi)\, E\, q^* \;\geq\; 0.$$

This is the politically-exploited share of structural opacity, and is the object that the Treasury disciplines.

The Treasury's utility function is

$$U_T \;=\; -\lambda \left[\Delta_{FI}(N, \xi)\right]^2 \;-\; \rho \left[\phi_0 + \phi_\xi\, \xi\right], \qquad \lambda, \rho, \phi_0, \phi_\xi > 0, \tag{8}$$

where $\lambda > 0$ measures the Treasury's preference for fiscal transparency (its institutional strength), $\rho > 0$ measures the cost of being removed from office, and $P(\text{removal}) = \phi_0 + \phi_\xi\, \xi$ is the removal probability, strictly increasing in the Treasury's enforcement effort. The first term penalizes the quadratic fiscal illusion gap: a Treasury endowed with strong institutional weight λ values transparency and dislikes a fiscal illusion gap. The second term captures the political cost of enforcement: the politician retaliates against a Treasury that aggressively exposes the hidden portion of tax expenditures, internalizing the political cost of enforcement and ensuring an interior $\xi^* \in (0, 1)$ without need of an ad-hoc cost-of-enforcement term.

Budget transparency and the rule of law. The institutional weight λ in the quadratic loss $-\lambda\,[\Delta_{FI}]^2$ determines whether budget transparency is a binding objective or a discretionary target. Its institutional foundations lie in the legal architecture of public finance (constitutional balanced-budget rules, fiscal responsibility legislation, medium-term fiscal frameworks, independent fiscal councils with legal mandates, and the credibility of the rule of law in public finance more broadly). We therefore interpret λ as the model-theoretic image of the rule of law specialized to fiscal discipline, which is the structural foundation for using the Rule of Law Index alongside social-capital measures as empirical proxies of Treasury enforcement strength in Section 5.

Equilibrium enforcement. Anticipating the politician's response, the Treasury's first-order condition (derived in Appendix A.3) delivers an implicit equation that pins down ξ^* as a function of the institutional ratio λ/ρ and of the equilibrium values of the remaining variables. The substantively relevant property is the strict monotonicity

$$\xi^* \text{ is strictly increasing in } \lambda/\rho. \tag{9}$$

A higher λ raises the marginal disutility from the fiscal illusion gap and pushes ξ^* up; a higher ρ raises the marginal disutility from removal and pushes ξ^* down. The ratio λ/ρ therefore summarizes the Treasury's net institutional incentive to enforce, and is the substantive property that drives all comparative statics that follow.

4.5. Equilibrium and Market Clearing

We solve for a subgame-perfect Nash equilibrium in which (i) at Stage 1 the firm chooses q and demands E_d given anticipated political supply, (ii) at Stage 2 the Treasury commits to enforcement level ξ, and (iii) at Stage 3 the politician chooses (G, E, N) subject to the enforced perception, with market clearing requiring $E = E_d$ at Stage 1. Backward induction implies that the politician takes the equilibrium output q^* pinned down at Stage 1 as parametric when choosing (G, E, N). The market-clearing condition equates the politician's supply of tax expenditures with the firm's demand:

$$\frac{\beta - \gamma\,\Omega(N^*, \xi^*)\,q^*}{k_E} \;=\; \frac{\theta(A-c)}{2b\mu - \theta^2}, \tag{10}$$

with q^* given by (3). The system formed by (6), (9) and (10) pins down the equilibrium values (G^*, E^*, N^*, ξ^*). The full backward-induction derivation, including the politician's Stage 3 problem with the explicit choice of N, the Treasury's Stage 2 first-order condition with the induced response of E via the implicit function theorem, and the existence argument, is collected in Appendix A.4.

4.6. Main Results

The solution of the model delivers a proposition and a corollary that connect institutional strength to the equilibrium composition of fiscal policy. The proposition establishes the comparative statics of (ξ^*, E^*, N^*, G^*) with respect to the Treasury's institutional ratio λ/ρ; the corollary specializes the result to the empirical asymmetry between the fragmentation margin and the volume margin.

Proposition 4.1 (Treasury Strength, Volume and Fragmentation)**.** Under the maintained assumptions on the model primitives and under the small-k_T approximation, the equilibrium (E^*, N^*, G^*, ξ^*) satisfies the following comparative statics with respect to λ/ρ:

1. $\partial\xi^*/\partial(\lambda/\rho) > 0$: a stronger Treasury enforces more in equilibrium;
2. $\partial E^*/\partial(\lambda/\rho) < 0$: stronger enforcement reduces the equilibrium level of tax expenditures;
3. $\partial N^*/\partial(\lambda/\rho) < 0$: stronger enforcement reduces the equilibrium number of distinct provisions;

4. $\partial G^*/\partial(\lambda/\rho) \approx 0$: visible public goods are approximately independent of Treasury strength under the small-k_T approximation.

The proof, collected in Appendix A.5, establishes Part (i) directly from the Treasury's first-order condition derived in Section 4.4, then establishes Parts (ii) and (iii) jointly via the implicit function theorem applied to the politician's first-order conditions in (E, N) taking ξ^* as parametric, with the determinant of the Jacobian positive by the second-order condition of the politician's problem. The intuition behind each result is direct: (i) a higher λ/ρ raises the Treasury's net incentive to enforce; (ii) the higher ξ^* raises Ω, which lowers E^* through the politician's first-order condition; (iii) the lower E^* and the lower $(1-\xi^*)$ both reduce the marginal political benefit of an additional provision in the politician's first-order condition for N, forcing N^* down because $\psi(\cdot)$ is strictly convex; (iv) G^* is approximately independent of (N, ξ) under the small-k_T approximation and is therefore not the relevant margin of adjustment.

What the model predicts about volume. Within the model as formally specified, Part (ii) of Proposition 4.1 predicts that an empirical proxy of λ/ρ should be *negatively* associated with E^*. Aggregate revenue foregone, $E^* \cdot q^*$, inherits this sign because q^* is independent of λ/ρ.

Corollary 4.2 (The Volume-Fragmentation Asymmetry)**.** The empirical counterpart of N^* admits no source of variation independent of the opacity mechanism formalized in this paper: fragmentation is generated only through the politician's choice of N governed by the first-order condition $\psi'(N^*) = \gamma\,\delta'(N^*)(1-\xi^*)E^*\,q^*$, so N^* is a structurally clean signal of opacity-driven capture. By contrast, the empirical counterpart of aggregate revenue foregone $E\,q^*$ is generated, in addition to the opacity-driven mechanism captured by the present model, by a legitimate *redistributive demand* component that is shaped by the size of the welfare state, demographic structure, and the policy goals delegated to the tax system (e.g., earned-income credits, R&D incentives). This redistributive component is not formally part of the model; an explicit extension nesting a redistribution motive in the politician's utility function is left to future work. As a result, the projection of empirically observed revenue foregone on λ/ρ produces a coefficient whose magnitude is attenuated by the redistributive component and whose sign may be empirically ambiguous, conditional on the controls included in the regression.

Empirical implication: the count of tax expenditure provisions (N) is a structurally cleaner empirical signal of opaque political capture than the aggregate fiscal magnitude ($E\,q^*$), because N^* has no plausible source of variation independent of the opacity mechanism, whereas $E\,q^*$ admits substantial variation driven by legitimate redistributive demand. Cross-country regressions of institutional quality on N should therefore deliver larger (in absolute value) and more statistically robust coefficients than analogous regressions on revenue foregone, even though both quantities are theoretically signed negatively in the absence of the redistributive component.

Numerical illustration. Figure 3 provides a numerical illustration of the joint comparative statics of E/G over (δ, λ) under a specific parametrisation. The bivariate surface shows that the response of E/G to λ is sharper at high values of δ and more gradual at intermediate

values, illustrating the joint role of opacity and enforcement in shaping the equilibrium composition of fiscal policy. Under the small-k_T approximation, the variation in E/G along λ is quantitatively driven by the contraction of E^* in the numerator, with G^* in the denominator approximately constant. The figure is intended as a numerical illustration of the comparative statics rather than as a formal characterization of a threshold. A one-dimensional slice of the surface, at fixed $\delta = 0.95$, is reported in Appendix B for readers interested in the marginal effect of λ in isolation.

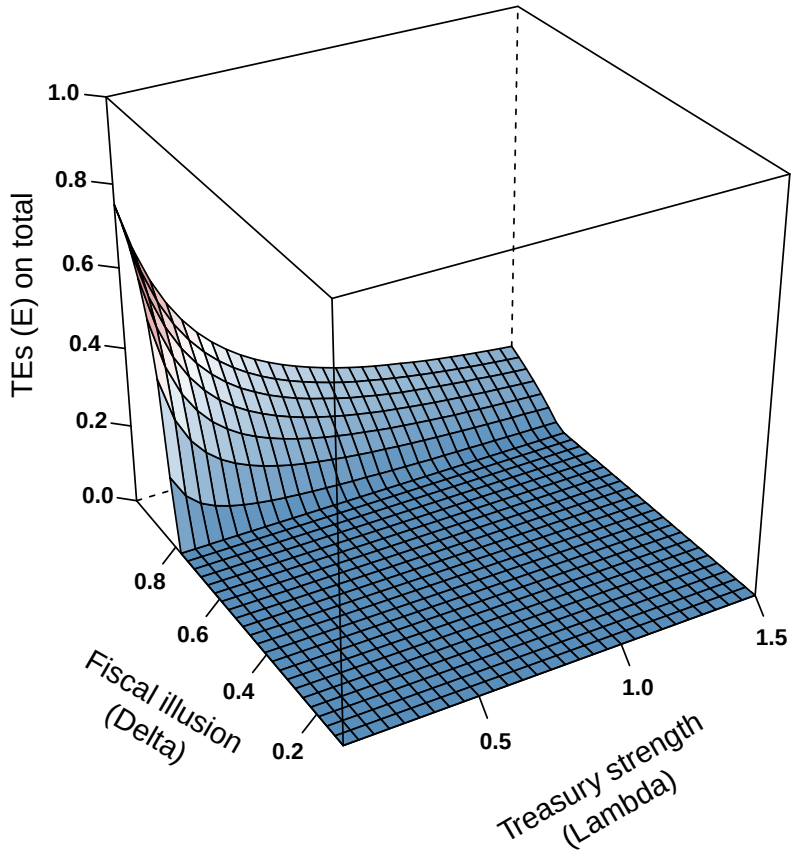

Figure 3: Numerical illustration of the joint comparative statics over (δ, λ). The response of E/G to λ is sharper at high values of δ and more gradual at intermediate values. A sufficiently strong Treasury ($\lambda > 1$) is associated with low equilibrium reliance on opaque tax expenditures across the range of δ values considered. The figure is illustrative of the model's qualitative comparative statics under a specific parametrisation.

Three features of the surface are worth highlighting. First, in the upper-right region of the parameter space (high structural opacity, $\delta \approx 0.9$, and weak institutional enforcement, $\lambda \approx 0.1$), the equilibrium share of tax expenditures is high: the politician faces a low effective political cost of opaque redistribution and provides tax expenditures at levels approaching the maximum of the parameter range. Second, the response of the equilibrium share to λ is non-linear: at high δ, a moderate increase in λ is associated with a substantial reduction in the share of tax expenditures, whereas at lower δ the same increase has more muted effects. Third, when the Treasury is institutionally strong ($\lambda > 1$), the equilibrium share remains low across the range of δ values considered. The policy reading suggested by the

numerical exercise is that strengthening the budget authority can reduce reliance on opaque tax expenditures even where the latent fiscal illusion is structurally high. The reading is illustrative, not a formal welfare result.

Welfare considerations. The fragmentation mechanism generates three distortions with welfare implications. First, latent fiscal illusion biases political choices toward hidden redistribution even where visible public goods would be more efficient. Second, rent-seeking competition among interest groups dissipates resources. Third, the diffusion of opacity across many narrow provisions weakens fiscal discipline through the proliferation-driven mechanism formalized in the politician's first-order condition for N. A formal welfare derivation requires explicit modelling of the social value of visible public goods, which is beyond the scope of the present paper; we therefore treat the welfare implications as suggestive and concentrate the rest of the analysis on the testable fragmentation channel.

5. Testable Hypotheses

The model is built around two outcomes that are observed separately, the number of provisions N and the aggregate fiscal magnitude $E\,q^*$, and around three structural parameters with natural empirical counterparts: the Treasury's institutional weight λ, whose foundations lie both in the legal architecture of budget transparency and in the social monitoring of the budget exercised by a civic society with high generalized trust; the rent-seeking cost μ, inversely related to corruption; and visible public goods G. Proposition 4.1 and Corollary 4.2 deliver three hypotheses, to be tested in cross-country data in which provisions are counted separately from revenue foregone.

Hypothesis 5.1 (Institutional strength reduces fragmentation)**.** Since a stronger Treasury, and the social monitoring that supports it, reduces the share of structural opacity δ that is actually exploited, it raises the political cost of opaque redistribution. Politicians respond by reducing the *number* of targeted provisions, as fragmentation maximises opacity. Empirical proxies of λ/ρ should therefore be negatively associated with the count of distinct tax expenditure provisions [Proposition 4.1(iii)], and proxies of cheap political influence (low μ) positively associated with it [equation (3)].

Hypothesis 5.2 (Institutional strength induces substitution from E to G)**.** Stronger institutional enforcement reduces the equilibrium level of tax expenditures while leaving visible public goods approximately unchanged [Proposition 4.1(ii) and (iv)]. It should therefore be associated with a shift in the composition of fiscal policy away from opaque tax expenditures and towards visible public goods, that is, with a lower ratio of tax expenditures to government spending.

Hypothesis 5.3 (Political capture operates through complexity)**.** Since the structural opacity δ is a function of the number of provisions, $\delta = \delta(N)$, the fragmentation count is the *direct* empirical counterpart of the latent opacity parameter, while the aggregate volume of revenue foregone mixes opacity-driven with legitimate redistributive components [Corollary 4.2]. Institutional quality should therefore constrain *fragmentation* robustly, whereas its effect on *aggregate volume* is expected to be muted, noisy or sign-ambiguous: the coefficient of institutional proxies on the count of provisions should be larger in absolute value and more robust than the coefficient on revenue foregone.

Two features of these hypotheses matter for their empirical implementation. First, Hypothesis 5.3 is a statement about the relative behavior of two outcomes under the same source of variation in institutional quality: whatever identification strategy is adopted for λ, whose endogeneity to fiscal practice is the main obstacle, the model predicts that it should bite on the count of provisions more clearly than on the headline cost. Second, since reporting conventions may affect the count of provisions, the quality of official tax-expenditure reporting is a necessary control. Taking the three hypotheses to the data is the natural next step of this research.

6. Concluding Remarks

Tax expenditures have two dimensions, their cost and their fragmentation into distinct provisions, and the literature has measured, regulated and modeled the first while treating the second as a descriptive curiosity. This paper has argued that the order should be reversed. Fragmentation is the channel through which the opacity of tax-based redistribution is produced, and once it is modeled as a choice the political economy of tax expenditures acquires a structure that the headline cost cannot reveal: the number of provisions is a structurally clean signal of institutionally disciplined capture, whereas aggregate revenue foregone, which mixes the opacity-driven component with legitimate redistributive demand, responds to institutions in an attenuated and possibly ambiguous way.

The policy message is that curbing tax expenditures is an institutional problem before it is a fiscal one. What disciplines hidden redistribution is a Treasury with a strong institutional weight on transparency and a low political cost of enforcing it, which is what independent fiscal councils with a mandate on tax expenditures, reporting obligations fixed in law and the budgetary unification of tax expenditures with spending are designed to provide; the numerical illustration suggests that strengthening the budget authority can reduce reliance on opaque tax expenditures even where the latent fiscal illusion is structurally high.

Several questions remain open: a formal welfare analysis with an explicit social value of visible public goods, an extension nesting a redistributive motive in the politician's utility, and the dynamics of fragmentation as provisions accumulate over successive legislatures. On the empirical side, the central prediction, that institutional quality disciplines the fragmentation of tax expenditures much more clearly than their volume, is sharp enough to be rejected wherever provisions are counted separately from revenue foregone, and taking it to the data is the next step.

References

Aghion, P., Algan, Y., Cahuc, P., Shleifer, A., 2010. Regulation and distrust. The Quarterly Journal of Economics 125, 1015–1049.

Alesina, A., Perotti, R., 1996. Fiscal discipline and the budget process. American Economic Review 86, 401–407.

Alesina, A., Tabellini, G., 2007. Bureaucrats or politicians? Part I: A single policy task. American Economic Review 97, 169–179.

Alt, J.E., Lassen, D.D., 2006. Fiscal transparency, political parties, and debt in oecd countries. European Economic Review 50, 1403–1439.

Becker, G.S., 1983. A theory of competition among pressure groups for political influence. The Quarterly Journal of Economics 98, 371–400.

Bordalo, P., Gennaioli, N., Shleifer, A., 2013. Salience and consumer choice. Journal of Political Economy 121, 803–843. doi:10.1086/673885.

Boycko, M., Shleifer, A., Vishny, R.W., 1994. Voucher privatization. Journal of Financial Economics 35, 249–266.

Boycko, M., Shleifer, A., Vishny, R.W., 1996. A theory of privatisation. Economic Journal 106, 309–319.

Buchanan, J.M., Wagner, R.E., 1980. Democracy in Deficit: The Political Legacy of Lord Keynes. Liberty Fund, Indianapolis.

Burman, L.E., Geissler, C., Toder, E.J., 2008. How big are total individual income tax expenditures, and who benefits from them? American Economic Review 98, 79–83.

Burman, L.E., Phaup, M., 2012. Tax expenditures, the size and efficiency of government, and implications for budget reform. Tax Policy and the Economy 26, 93–124.

Chetty, R., Looney, A., Kroft, K., 2009. Salience and taxation: Theory and evidence. American Economic Review 99, 1145–1177.

Debrun, X., Kinda, T., 2017. Strengthening post-crisis fiscal credibility: Fiscal councils on the rise – a new dataset. Fiscal Studies 38, 667–700.

Finkelstein, A., 2009. E-ztax: Tax salience and tax rates. The Quarterly Journal of Economics 124, 969–1010.

Gabaix, X., 2014. A sparsity-based model of bounded rationality. Quarterly Journal of Economics 129, 1661–1710.

Grossman, G.M., Helpman, E., 1994. Protection for sale. American Economic Review 84, 833–850.

Grossman, G.M., Helpman, E., 2001. Special Interest Politics. MIT Press, Cambridge, MA.

Guiso, L., Sapienza, P., Zingales, L., 2011. Civic capital as the missing link, in: Benhabib, J., Bisin, A., Jackson, M.O. (Eds.), Handbook of Social Economics. North-Holland. volume 1, pp. 417–480.

von Hagen, J., Harden, I.J., 1995. Budget processes and commitment to fiscal discipline. European Economic Review 39, 771–779.

Hallerberg, M., Strauch, R., Von Hagen, J., 2004. The design of fiscal rules and forms of governance in european union countries. European Journal of Political Economy 20, 563–589.

Kleinbard, E.D., 2010. The congress within the congress: How tax expenditures distort our budget and our political processes. Ohio Northern University Law Review 36, 1–30.

Krueger, A.O., 1974. The political economy of the rent-seeking society. American Economic Review 64, 291–303.

La Porta, R., Lopez-de Silanes, F., Shleifer, A., Vishny, R.W., 1997. Trust in large organizations. American Economic Review Papers and Proceedings 87, 333–338.

Milesi-Ferretti, G.M., 2004. Good, bad or ugly? On the effects of fiscal rules with creative accounting. Journal of Public Economics 88, 377–394.

Niskanen, W.A., 1971. Bureaucracy and Representative Government. Aldine-Atherton, Chicago.

Oates, W.E., 1988. On the nature and measurement of fiscal illusion: A survey, in: Brennan, G., et al. (Eds.), Taxation and Fiscal Federalism: Essays in Honour of Russell Mathews. Australian National University Press, pp. 65–82.

Persson, T., Tabellini, G., 2000. Political Economics: Explaining Economic Policy. MIT Press, Cambridge, MA.

Puviani, A., 1903. Teoria della Illusione Finanziaria. Remo Sandron, Palermo.

Redonda, A., von Haldenwang, C., Aliu, F., 2023. The global tax expenditures database (gted). Journal of International Development 35, 2543–2560.

Rogoff, K., 1985. The optimal degree of commitment to an intermediate monetary target. Quarterly Journal of Economics 100, 1169–1189.

Shleifer, A., Vishny, R.W., 1994. Politicians and firms. The Quarterly Journal of Economics 109, 995–1025.

Surrey, S.S., 1973. Pathways to Tax Reform: The Concept of Tax Expenditures. Harvard University Press, Cambridge, MA.

Tabellini, G., 2008. Institutions and culture. Journal of the European Economic Association 6, 255–294.

Tax Administration Research Centre, 2014. The Definition, Measurement, and Evaluation of Tax Expenditures and Tax Reliefs. Research Report for the National Audit Office. Tax Administration Research Centre, University of Exeter. Exeter.

Tullock, G., 1967. The welfare costs of tariffs, monopolies, and theft. Western Economic Journal 5, 224–232.

Tyson, S.A., 2014. The distinct political economy of tax expenditures. American Political Science Review 108, 788–801.

Villela, L., Lemgruber, A., Jorratt, M., 2010. Tax Expenditure Budgets: Concepts and Challenges for Implementation. IDB Working Paper Series IDB-WP-131. Inter-American Development Bank. Washington, DC.

Wyplosz, C., 2013. Fiscal rules: Theoretical issues and historical experiences, in: Alesina, A., Giavazzi, F. (Eds.), Fiscal Policy after the Financial Crisis. University of Chicago Press, pp. 495–525.

Appendix A. Theoretical model: derivations and proofs

This appendix collects the formal derivations and proofs of the three-agent model presented in Section 4. The structure mirrors the order of the main text: Section Appendix A.1 derives the firm's equilibrium output and demand for tax expenditures; Section Appendix A.2 derives the politician's first-order conditions under the small-k_T approximation, with the explicit role of the quadratic parametrisation $\psi(N) = (\psi_0/2)N^2$; Section Appendix A.3 derives the Treasury's enforcement choice including the induced response of E via the implicit function theorem; Section Appendix A.4 ties the three stages together; Section Appendix A.5 proves Proposition 4.1.

Appendix A.1. Derivation of the firm's optimum

The firm's problem stated in equation (2) of Section 4.2 is

$$\max_{q,E} \ \pi \ = \ (A - bq)q - (c - \theta E)q - \tfrac{\mu}{2}E^2.$$

Taking first-order conditions with respect to q and E yields

$$\frac{\partial \pi}{\partial q} = A - 2bq - c + \theta E = 0, \tag{Appendix A.1}$$

$$\frac{\partial \pi}{\partial E} = \theta q - \mu E = 0. \tag{Appendix A.2}$$

Equation (Appendix A.2) delivers the firm's optimal demand for tax expenditures,

$$E_d \ = \ \frac{\theta q}{\mu}, \tag{Appendix A.3}$$

which, substituted into (Appendix A.1), yields

$$A - 2bq - c + \theta \cdot \frac{\theta q}{\mu} \ = \ 0 \quad \Longrightarrow \quad q^* \ = \ \frac{\mu(A-c)}{2b\mu - \theta^2}.$$

Substituting q^* back into (Appendix A.3) gives the equilibrium demand for tax expenditures:

$$E_d^* \ = \ \frac{\theta}{\mu} \cdot \frac{\mu(A-c)}{2b\mu - \theta^2} \ = \ \frac{\theta(A-c)}{2b\mu - \theta^2}. \tag{Appendix A.4}$$

The denominator is positive under the stability condition (4), $2b\mu > \theta^2$.

Appendix A.2. Derivation of the politician's optimum and the small-k_T approximation

The politician's perceived tax burden, given the Treasury's enforcement choice ξ and the effective visibility $\Omega(N, \xi) = 1 - \delta(N)(1 - \xi)$ defined in (7), is

$$T \ = \ G + \Omega(N, \xi)\, E\, q^*. \tag{Appendix A.5}$$

Substituting (Appendix A.5) into the politician's utility function (5) delivers the unconstrained problem

$$\max_{G,E,N} \ V \ = \ \alpha G - \tfrac{1}{2}k_G G^2 + \beta E - \tfrac{1}{2}k_E E^2 - \gamma\big[G + \Omega(N,\xi)\, E\, q^*\big] - \tfrac{1}{2}k_T\big[G + \Omega(N,\xi)\, E\, q^*\big]^2 - \psi(N).$$

The first-order conditions are

$$\frac{\partial V}{\partial G} = \alpha - k_G G - \gamma - k_T\big[G + \Omega E q^*\big] = 0, \qquad \text{(Appendix A.6)}$$

$$\frac{\partial V}{\partial E} = \beta - k_E E - \gamma\,\Omega\,q^* - k_T\big[G + \Omega E q^*\big]\,\Omega\,q^* = 0, \qquad \text{(Appendix A.7)}$$

$$\frac{\partial V}{\partial N} = -\gamma\,\tfrac{\partial \Omega}{\partial N}\,E\,q^* - k_T\big[G + \Omega E q^*\big]\,\tfrac{\partial \Omega}{\partial N}\,E\,q^* - \psi'(N) = 0. \qquad \text{(Appendix A.8)}$$

Since $\partial\Omega/\partial N = -\delta'(N)(1-\xi) < 0$, equation (Appendix A.8) can be rewritten as

$$\gamma\,\delta'(N)\,(1-\xi)\,E\,q^* + k_T\big[G + \Omega E q^*\big]\,\delta'(N)\,(1-\xi)\,E\,q^* \;=\; \psi'(N).$$

Under the small-k_T approximation ($k_T \to 0$), the second term on the left-hand side vanishes and the three first-order conditions collapse to

$$G^* \;\approx\; \frac{\alpha - \gamma}{k_G}, \qquad E^* \;\approx\; \frac{\beta - \gamma\,\Omega(N^*, \xi^*)\,q^*}{k_E}, \qquad \psi'(N^*) \;=\; \gamma\,\delta'(N^*)(1-\xi^*)\,E^*\,q^*,$$

which are the expressions reported in equation (6) of the main text. The full expression for G^*, retaining the k_T terms, is

$$G^* \;=\; \frac{\alpha - \gamma - k_T\,\Omega E q^*}{k_G + k_T}, \qquad \text{(Appendix A.9)}$$

which converges to $(\alpha - \gamma)/k_G$ as $k_T \to 0$.

Quadratic parametrisation and implicit form for N^*. Adopting the tractable parametrisation $\psi(N) = (\psi_0/2)N^2$ with $\psi_0 > 0$ delivers $\psi'(N) = \psi_0\,N$, and the first-order condition for N becomes

$$N^* \;=\; \frac{\gamma\,\delta'(N^*)\,(1-\xi^*)\,E^*\,q^*}{\psi_0}. \qquad \text{(Appendix A.10)}$$

With $\delta(N)$ specified as in (1), $\delta'(N) = \delta_{\max}\,\nu/(1+\nu N)^2$, so (Appendix A.10) reduces to the fixed-point equation

$$N^*\,(1+\nu N^*)^2 \;=\; \frac{\gamma\,\delta_{\max}\,\nu\,(1-\xi^*)\,E^*\,q^*}{\psi_0}.$$

The left-hand side is a strictly increasing function of N^* from 0 to ∞, so the equation has a unique positive solution for any value of the right-hand side. Differentiating implicitly with respect to ξ^* delivers $\partial N^*/\partial \xi^* < 0$, confirming the qualitative property used in the proof of Proposition 4.1. We maintain the general convex $\psi(\cdot)$ in what follows, since the substantive results do not rely on the quadratic form.

Appendix A.3. Derivation of the Treasury's enforcement choice

The Treasury anticipates the politician's Stage 3 response and chooses ξ to maximize (8). Recalling that $\Delta_{FI}(N, \xi) = \delta(N)(1-\xi)\,E\,q^*$, the first-order condition is

$$\frac{\partial U_T}{\partial \xi} \;=\; -2\lambda\,\Delta_{FI}(N^*, \xi)\,\frac{d\Delta_{FI}}{d\xi} \;-\; \rho\,\phi_\xi \;=\; 0, \qquad \text{(Appendix A.11)}$$

where the total derivative $d\Delta_{FI}/d\xi$ accounts for both the direct effect of ξ on Δ_{FI} and the indirect effect through the politician's best response $E^*(N^*, \xi)$.

From the politician's first-order condition (Appendix A.7), the implicit function theorem applied under the small-k_T approximation delivers

$$\left.\frac{\partial E^*}{\partial \xi}\right|_{N^*} \;=\; -\frac{\gamma\, q^*}{k_E}\,\frac{\partial \Omega}{\partial \xi} \;=\; -\frac{\gamma\, \delta(N^*)\, q^*}{k_E} \;<\; 0, \qquad \text{(Appendix A.12)}$$

since $\partial\Omega/\partial\xi = \delta(N^*) > 0$. The total derivative of Δ_{FI} with respect to ξ, holding N^* at its equilibrium value, is then

$$\begin{aligned}\left.\frac{d\Delta_{FI}}{d\xi}\right|_{N^*} &= -\delta(N^*)\, E^*\, q^* + \delta(N^*)(1-\xi)\, q^* \cdot \left[-\frac{\gamma\, \delta(N^*)\, q^*}{k_E}\right] \\ &= -\delta(N^*)\, q^* \left[E^* + \frac{\gamma\, \delta(N^*)\,(1-\xi)\, q^*}{k_E}\right] \;<\; 0.\end{aligned}$$

Substituting into (Appendix A.11) yields

$$2\lambda\, \delta(N^*)^2\,(1-\xi^*)\,(E^*)^2\,(q^*)^2 \left[1 + \frac{\gamma\, \delta(N^*)\,(1-\xi^*)\, q^*}{k_E\, E^*}\right] \;=\; \rho\, \phi_\xi. \qquad \text{(Appendix A.13)}$$

The left-hand side, denoted $H(\xi^*;\, \lambda, \rho, \ldots)$, is strictly decreasing in ξ^* on a neighbourhood of the interior solution: each constituent factor, $(1-\xi^*)$, $E^*(N^*, \xi^*)$, and the bracketed term, is non-increasing in ξ^*, and at least one strictly decreasing. Equation (Appendix A.13) therefore admits a unique interior solution $\xi^* \in (0, 1)$ whenever $H(0) > \rho\, \phi_\xi$, which is the maintained existence condition.

Monotonicity. The implicit equation (Appendix A.13) implies $\partial\xi^*/\partial\lambda > 0$ and $\partial\xi^*/\partial\rho < 0$, because H is independent of (λ, ρ), the left-hand side is decreasing in ξ^*, and the right-hand side $\rho\, \phi_\xi$ is increasing in ρ. Combining the two signs delivers

$$\frac{\partial \xi^*}{\partial(\lambda/\rho)} \;>\; 0,$$

which is the property stated in (9) and used throughout the comparative-statics analysis.

Appendix A.4. Market clearing and equilibrium existence

The market-clearing condition at Stage 1 equates the politician's supply of tax expenditures with the firm's demand:

$$E^* \;=\; E_d^* \;\Longleftrightarrow\; \frac{\beta - \gamma\, \Omega(N^*, \xi^*)\, q^*}{k_E} \;=\; \frac{\theta(A-c)}{2b\mu - \theta^2},$$

which reproduces equation (10) of the main text. With q^* given by the expression derived in Section Appendix A.1, the system formed by (6), (9) and (10), augmented by the first-order condition for N^* derived in Section Appendix A.2, pins down the equilibrium values $(G^*, E^*, N^*, \xi^*, q^*)$. Existence and uniqueness in a neighbourhood of the interior solution follow from the joint concavity of the politician's problem (guaranteed by $k_G, k_E > 0$ and by the strict convexity of ψ), the strict monotonicity of the function $H(\cdot)$ in ξ established in Section Appendix A.3, and a standard fixed-point argument on the joint system.

Appendix A.5. Proof of Proposition 4.1

We maintain throughout the small-k_T approximation introduced in Section Appendix A.2. The proof proceeds in three steps: Part (i) is established directly from Section Appendix A.3; Parts (ii) and (iii) are established *jointly* via the implicit function theorem applied to the politician's first-order conditions in (E, N), taking ξ^* as parametric; Part (iv) follows directly from (6).

Part (i): $\partial\xi^*/\partial(\lambda/\rho) > 0$. Established in Section Appendix A.3, equation (9).

Joint analysis of Parts (ii) and (iii) via implicit function theorem. Consider the politician's first-order conditions for (E, N) taking ξ as parametric:

$$\begin{aligned} F_1(E, N; \xi) &\equiv \beta - k_E\, E - \gamma\, \Omega(N, \xi)\, q^* \;=\; 0, \\ F_2(E, N; \xi) &\equiv \gamma\, \delta'(N)\, (1 - \xi)\, E\, q^* - \psi'(N) \;=\; 0. \end{aligned}$$

The Jacobian of (F_1, F_2) with respect to (E, N) coincides with the Hessian of the politician's payoff function:

$$J \;=\; \begin{pmatrix} -k_E & \gamma\, q^*\, \delta'(N^*)(1 - \xi^*) \\ \gamma\, q^*\, \delta'(N^*)(1 - \xi^*) & \gamma\, \delta''(N^*)(1 - \xi^*)\, E^*\, q^* - \psi''(N^*) \end{pmatrix}.$$

The second-order condition for an interior maximum requires $\det J > 0$, i.e.

$$k_E \left[\psi''(N^*) - \gamma\, \delta''(N^*)(1 - \xi^*)\, E^*\, q^*\right] \;>\; \gamma^2\, q^{*2} \left[\delta'(N^*)(1 - \xi^*)\right]^2,$$

which is the maintained regularity condition on the model primitives. The bracket on the left-hand side is strictly positive ($\psi'' > 0$, $\delta'' < 0$), so the SOC is a meaningful restriction on the relative magnitudes of ψ'', δ' and δ''.

The partial derivatives of (F_1, F_2) with respect to ξ are $\partial F_1/\partial\xi = -\gamma\, q^*\, \delta(N^*) < 0$ and $\partial F_2/\partial\xi = -\gamma\, \delta'(N^*)\, E^*\, q^* < 0$. Applying the implicit function theorem,

$$\begin{aligned} \frac{\partial E^*}{\partial \xi} &= -\frac{\gamma\, q^*}{\det J} \Big[\delta(N^*) \left(\psi''(N^*) - \gamma\, \delta''(N^*)(1 - \xi^*)\, E^*\, q^*\right) + \gamma\, q^* \left[\delta'(N^*)(1 - \xi^*)\right]^2 E^*\Big], \\ \frac{\partial N^*}{\partial \xi} &= -\frac{\gamma\, q^*\, \delta'(N^*)}{\det J} \Big[k_E\, E^* + \gamma\, q^*\, \delta(N^*)(1 - \xi^*)\Big]. \end{aligned}$$

Both numerators in square brackets are strictly positive (each term is positive: $\delta > 0$, $\psi'' > 0$, $-\delta'' > 0$, $E^* > 0$, $\delta' > 0$, $k_E > 0$). With $\det J > 0$ by the SOC, the leading minus sign delivers $\partial E^*/\partial\xi < 0$ and $\partial N^*/\partial\xi < 0$. By the chain rule with Part (i),

$$\frac{\partial E^*}{\partial(\lambda/\rho)} \;=\; \frac{\partial E^*}{\partial \xi} \cdot \frac{\partial \xi^*}{\partial(\lambda/\rho)} \;<\; 0, \qquad \frac{\partial N^*}{\partial(\lambda/\rho)} \;=\; \frac{\partial N^*}{\partial \xi} \cdot \frac{\partial \xi^*}{\partial(\lambda/\rho)} \;<\; 0,$$

which establishes Parts (ii) and (iii).

Part (iv): $\partial G^*/\partial(\lambda/\rho) \approx 0$. From (6), $G^* \approx (\alpha - \gamma)/k_G$, which is independent of (N^*, ξ^*) and hence of λ/ρ. The full expression (Appendix A.9) contains a correction of order k_T that vanishes under the maintained approximation. □

Appendix B. Numerical illustration: one-dimensional slice

Figure B.4 reports a one-dimensional slice of the bivariate surface shown in Figure 3 of the main text, holding the structural opacity at $\delta = 0.95$ to isolate the role of the Treasury enforcement parameter λ.

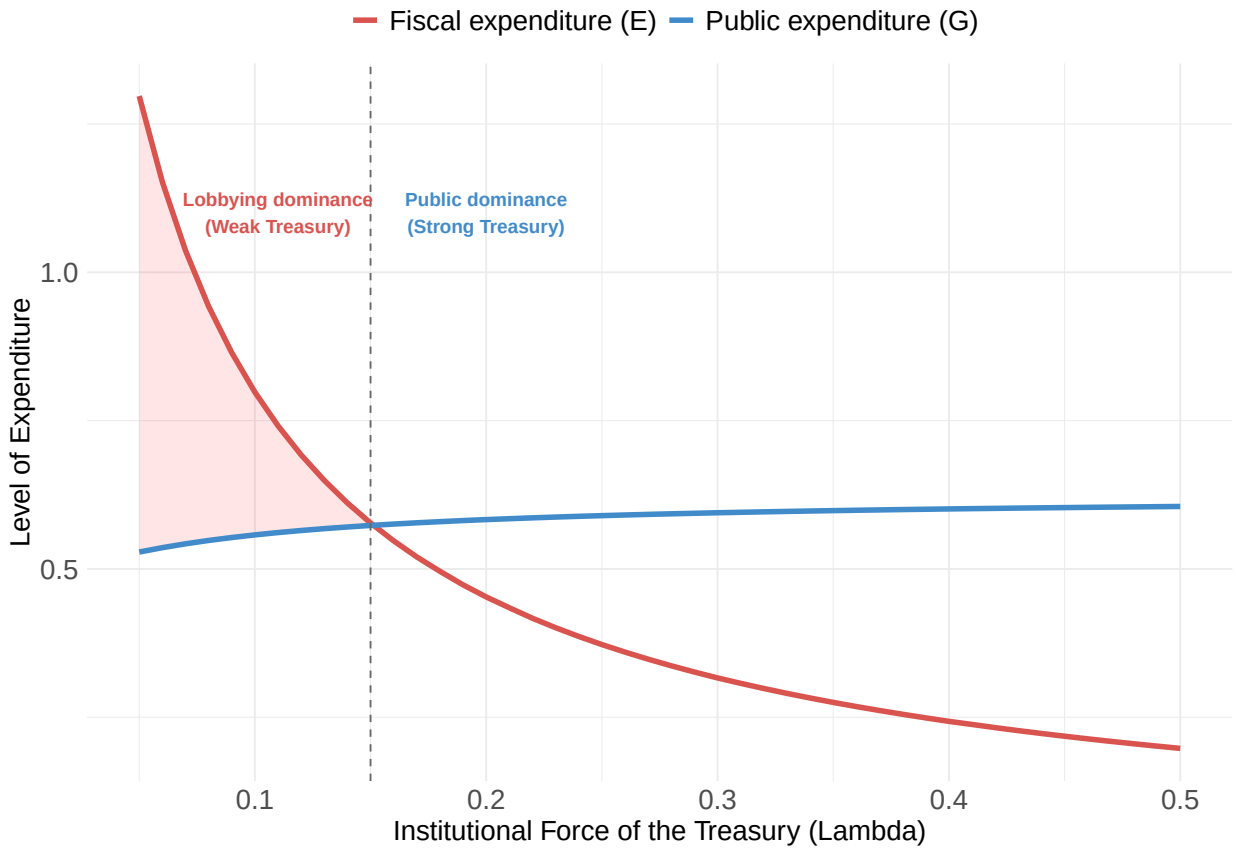

Figure B.4: Numerical illustration of the comparative statics in Proposition 4.1. Substitution between tax expenditures and visible public goods as the Treasury enforcement parameter λ varies, at fixed structural opacity $\delta = 0.95$, under a specific parametrisation. The plotted slice corresponds to $\delta = 0.95$ of the bivariate surface in Figure 3. The figure is illustrative of the negative comparative-static effect of λ on E/G and is not intended as a formal proof of a discontinuous threshold.

On the left side of the graph (the Weak Treasury domain), the equilibrium is dominated by rent-seeking. The high structural opacity of tax expenditures ($\delta \approx 0.95$) reduces the perceived political cost of tax expenditures, leading the politician to provide them at a high level to satisfy the lobbying firm. Visible public goods (G), being fully salient and politically costly to finance, are crowded out. As λ increases (the Strong Treasury domain), the increase in λ raises the effective visibility of tax expenditures and therefore their perceived political cost. The politician reduces tax expenditures and reallocates resources towards visible public goods, which become relatively more attractive once the budget constraint is more strictly enforced. The illustration shows that this transition can be highly non-linear under the given parametrisation, but the model as specified does not prove the existence of a true discontinuity, so the pattern is best read as a marked non-linear comparative static.