WORKING PAPERS

N° 1711

February 2026

"Optimal Taxation under Imperfect Trust"

Emin Ablyatifov and Georgy Lukyanov

Toulouse School of Economics

Optimal Taxation under Imperfect Trust

Emin Ablyatifov[1] and Georgy Lukyanov[2]

[1]Independent Researcher

[2]Toulouse School of Economics

Abstract

We study optimal taxation when the conversion of tax revenue into public goods is uncertain. In a static Ramsey framework with a representative household, a competitive firm, and two broad instruments (a labor-income tax and a commodity/output tax), a simple measure of trust— the perceived likelihood that revenue is actually delivered as public consumption—scales the marginal value of public funds. We show: (i) a trust threshold below which any distortionary taxation reduces welfare; (ii) above that threshold, policy uniquely pins down the scale of taxation but leaves a continuum of tax mixes (an equivalence frontier) that implement the same allocation and welfare; and (iii) tiny administrative or salience wedges select a unique instrument, typically favoring a broad base collected at source. We derive a trust-adjusted Ramsey rule in sufficient-statistics form, establish robustness to mild preference non-separabilities and concave public-good utility, and provide an isoelastic specialization with transparent comparative statics.

Keywords: Optimal taxation; public goods; credibility; marginal value of public funds; tax mix; administration.

JEL: E61, H21, H30, C73.

1 Introduction

Tax systems are built on a simple promise: citizens hand over resources today and receive public services in return. But in many places that promise is uncertain. Money leaks through waste, corruption, or weak capacity; delivery is delayed or diluted; and citizens doubt that an extra dollar of tax will show up as an extra dollar of roads, clinics, or security. When that doubt is material, the classic trade-off behind optimal taxation changes. This paper takes that doubt—call it trust in the conversion of revenue into public goods—and puts it at the center of a clean Ramsey benchmark.[1]

We study a static environment with a representative household, a competitive firm, and two broad tax instruments (a labor-income tax and a commodity or output tax).[2] Trust is treated as an exogenous, measurable belief about whether collected revenue will in fact be delivered as public consumption. Agents choose labor and consumption before the government's type is realized, so they internalize the expected conversion of taxes into public goods.[3]

Four results emerge. First, there is a sharp trust threshold:[4] below it, any distortionary taxation lowers welfare, and the unique optimum is no tax at all (Proposition 1). Second, once trust clears that threshold, policy pins down the overall scale of taxation but leaves the tax mix indeterminate along a one-dimensional equivalence frontier; every point on that frontier implements the same allocation and welfare (Theorem 1, with comparative statics in Proposition 3 and global concavity in Proposition 4). Third, tiny, instrument-specific frictions—administrative costs or salience differences—break the indifference and uniquely select the cheaper instrument, typically the broad base collected at source (Proposition 5). Fourth, we derive a trust-adjusted Ramsey rule in sufficient-statistics form: the marginal excess burden per marginal unit of delivered revenue equals a trust-scaled marginal value of public funds (Section 5.6).

Robustness is straightforward. Mild non-separability between consumption and leisure leaves the scale result intact but predictably tilts the preferred mix: when consumption and leisure are complements, the planner leans toward the commodity base; when they are substitutes, toward the labor base (Proposition 6). Allowing diminishing marginal utility of public goods reduces the optimal scale but preserves the frontier logic and the selection-by-wedges result; we characterize the scale implicitly and show it moves monotonically with trust (Propositions 7–8, Theorem 2). An isoelastic specialization provides closed-form schedules and figures that map measured trust into optimal rates, delivered public consumption, and welfare.

[1] We treat "trust" as a reduced-form, exogenous sufficient statistic for credibility, administrative capacity, and leakage risk. Keeping it exogenous in a one-period benchmark delivers closed forms and avoids overlap with the dynamic reputation study, where trust evolves endogenously.

[2] Think of the commodity/output tax as a uniform VAT or sales tax applied to a single composite good. In a representative-agent, one-good setting with no intermediate distortions, this aligns with broad-base recommendations from Diamond and Mirrlees (1971a,b); Atkinson and Stiglitz (1976). Multi-sector or input-tax complications are beyond our scope here.

[3] This timing implies an expected-utility evaluation of the public-good lottery. If agents were risk-averse over that lottery or if $v(\cdot)$ were concave, the trust threshold rises; we characterize this in Section 6.

[4] The threshold depends only on primitives at the origin (local marginal deadweight loss versus marginal delivered revenue). With concave public-good utility v, it becomes $1/\big(v'(0)Y^*\big)$; see Section 6.

2 Related Literature

Our benchmark is a static Ramsey problem with distortionary instruments financing a public good. The classic production-efficiency and broad-base results of Diamond and Mirrlees (1971a,b) and the uniform-tax logic of Atkinson and Stiglitz (1976) underpin our equivalence frontier: when trust is sufficiently high to warrant positive taxation, the planner is indifferent across tax mixes that keep the private distortion index fixed while raising the same delivered revenue. The sufficient-statistics view echoes the spirit of classic Ramsey calculations (Ramsey, 1927) while making the marginal value of public funds *trust-adjusted*.

The selection of a unique instrument on the frontier by tiny wedges relates to the administrative/evasion perspective in Slemrod and Yitzhaki (2002). Our prediction that the broad base is typically favored is consistent with institutional evidence and practice documented in Ebrill et al. (2001) and field evidence on VAT self-enforcement (Pomeranz, 2015). In our static model, such differences enter as curvature in instrument-specific costs and pin down the mix.

Time inconsistency and rules versus discretion (Kydland and Prescott, 1977) triggered the reputation literature in macroeconomic policy (Barro and Gordon, 1983; Barro, 1986) and a broader commitment/sustainable-plans agenda (Chari and Kehoe, 1990; Debortoli and Nunes, 2010). We take a complementary route: by *keeping trust exogenous* in a one-shot benchmark, we obtain (i) a sharp *trust threshold* for any taxation and (ii) a one-dimensional *equivalence frontier* in the tax mix. Dynamic reputation models with hidden types and updating—e.g. Phelan (2006) and Lu (2013)—are natural complements; our sufficient-statistics formulas provide primitives for such environments.

Empirically, higher institutional trust and tax morale correlate with stronger compliance and fiscal capacity, e.g. Scholz and Lubell (1998) and Torgler (2007). While our benchmark keeps compliance at one (and thus abstracts from evasion), the trust-scaling of the marginal value of public funds aligns with the broader finding that legitimacy conditions the effectiveness of taxation. Extensions with compliance responses can be added without leaving the static framework (see Section 6 for other robustness results).

Roadmap. Section 3 introduces the environment and instruments. Section 4 states the planner's problem and defines sufficient statistics. Section 5 presents the main results: the trust threshold, the equivalence frontier, instrument selection with small wedges, and the trust-adjusted Ramsey rule. Section 6 provides robustness: mild non-separability between consumption and labor, and a general concave utility for the public good. Section 7 specializes to an isoelastic case with closed forms and figures. Section 8 offers policy implications. Section

9 concludes. Proofs are collected in Appendix A.

3 Environment

Time is a single period. The private good is the numéraire. The economy has a representative household, a competitive firm, and a government.

The household has strictly concave $u(C, G, L)$ with $u_C > 0$, $u_G > 0$, $u_L < 0$, and standard curvature. We permit additivity in G as a leading case: $u(C, G, L) = \tilde{u}(C, L) + v(G)$ with $\tilde{u}_C > 0$, $\tilde{u}_L < 0$, and $v'(\cdot) > 0$, $v''(\cdot) \leq 0$. The household is atomistic and takes G as given when choosing L.

A competitive firm produces $Y = f(L)$ with $f' > 0$, $f'' < 0$ and pays wage w; profits Π accrue to the household. Let $k \equiv f'(L)L$ denote the private marginal product times input at the chosen L.

The government levies proportional taxes $\tau_\ell \in [0, 1)$ on labor income and $\tau_c \in [0, 1)$ on the broad commodity/output base.[5]

After private choices are made, the government's type realizes. With probability $\theta \in (0, 1)$ it is *honest* and converts collected revenue one-for-one into G; with probability $1 - \theta$ it is *opportunistic* and provides $G = 0$.[6]

In the honest realization, the public-goods budget identity is $G = \tau_\ell w L + \tau_c Y$.

The government announces (τ_ℓ, τ_c); private agents choose (C, L) and input demand; then the type realizes and, if honest, transforms revenue into G.

Given (τ_ℓ, τ_c), the firm solves $\max_{L \geq 0} (1 - \tau_c) f(L) - wL$, yielding $w = (1 - \tau_c) f'(L)$. Profits are $\Pi = (1 - \tau_c)[f(L) - f'(L)L]$. The household's budget constraint binds:

$$C = w(1 - \tau_\ell)L + \Pi. \tag{1}$$

Combining firm conditions and (1) implies (suppressing the argument L)

$$C = (1 - \tau_c)\,(1 - \alpha\,\tau_\ell)\,f(L) \quad \text{where} \quad \alpha \equiv \frac{f'(L)L}{f(L)} \in (0, 1). \tag{2}$$

For the isoelastic case $f(L) = aL^\beta$ we have $\alpha = \beta$ (constant). The household chooses L to maximize $u(C, G, L)$ taking G as given; in the additively separable benchmark this is

[5] We treat τ_c as a uniform tax on the single private good/value-added; with representative agents and no intermediate distortions, this corresponds to a broad commodity or output tax.

[6] This binary type can be read as a reduced-form for leakage, waste, or non-delivery. An equivalent interpretation is Bernoulli delivery of public goods with success probability θ; allowing fractional waste does not affect the main insights.

equivalent to maximizing $\tilde{u}(C(L), L)$. Existence and uniqueness of L^* follow from standard concavity assumptions. Let $Y^* \equiv f(L^*)$ and $k^* \equiv f'(L^*)L^*$.

Assumption 1. (i) $f : \mathbb{R}_+ \to \mathbb{R}_+$ is C^2, strictly concave, strictly increasing, with $f(0) = 0$ and $f(L) > 0$ for $L > 0$.
(ii) $\tilde{u} : \mathbb{R}_{++} \times \mathbb{R}_+ \to \mathbb{R}$ is C^2, strictly concave, nondecreasing in C and nonincreasing in L. In the benchmark we take $\tilde{u}(C, L) = \ln C - \phi(L)$ with ϕ strictly convex, $\lim_{L\to\infty} \phi(L) = +\infty$.
(iii) $v : \mathbb{R}_+ \to \mathbb{R}$ is C^2, nondecreasing and concave; the leading case is $v(G) = G$.
(iv) For any $(\tau_\ell, \tau_c) \in [0, 1)^2$, the household problem admits a unique interior solution $L^*(\tau_\ell, \tau_c) > 0$.

Assumption 2. Let $\alpha^* \equiv \frac{f'(L^*)L^*}{f(L^*)} \in (0, 1)$, $Y^* \equiv f(L^*)$, and $k^* \equiv f'(L^*)L^* = \alpha^* Y^*$. Define

$$S(\tau_\ell, \tau_c) \equiv (1 - \tau_c)\big(1 - \alpha^*\, \tau_\ell\big), \qquad \widetilde{R}(\tau_\ell, \tau_c) \equiv \tau_\ell(1 - \tau_c) + \frac{1}{\alpha^*}\, \tau_c.$$

Then at the induced private equilibrium,

$$C^* = S(\tau_\ell, \tau_c)\, Y^*, \qquad G^B(\tau_\ell, \tau_c) = k^*\, \widetilde{R}(\tau_\ell, \tau_c) \;=\; k^*\, \tau_\ell(1 - \tau_c) + Y^*\, \tau_c.$$

Remark 1. This normalization ensures the marginal delivered revenue effects match primitives: $\partial_{\tau_\ell} G^B|_{(0,0)} = k^*$ and $\partial_{\tau_c} G^B|_{(0,0)} = Y^*$.

4 Planner and Sufficient Statistics

This section recasts the policy problem in a form that makes the role of trust transparent and portable across primitives. Rather than work directly with tax rates, we map policies into two sufficient statistics:[7] a single "private distortion" index that captures how the tax system loads on the net-of-tax margins, and a delivered-revenue term that captures how much public consumption is expected when policy is implemented. This index mapping lets us express the planner's objective in a compact way and pin down what is identified by data or calibration (e.g., measured trust, administrative wedges) and what is a genuine choice variable (the tax mix versus the overall scale). With these building blocks, the main results in Section 5 follow from first principles and require only mild regularity.

A benevolent planner (ex ante) chooses (τ_ℓ, τ_c) to maximize expected utility

$$W(\tau_\ell, \tau_c; \theta) = \tilde{u}(C^*(\tau_\ell, \tau_c),\, L^*(\tau_\ell, \tau_c)) + \theta\, v\big(G^B(\tau_\ell, \tau_c)\big)\,, \tag{3}$$

[7] Our normalization for delivered revenue ensures that the marginal delivered revenue from a unit labor-tax increase equals k^* (the marginal product times input) and from a unit commodity-tax increase equals Y^* (output). This choice yields the clean threshold and frontier expressions used later.

where

$$G^B(\tau_\ell, \tau_c) = \tau_\ell\, w^* L^* + \tau_c\, Y^* = [\tau_\ell(1-\tau_c)\,\alpha^* + \tau_c]\, Y^*, \tag{4}$$

is delivered public consumption in the honest state (a function of (τ_ℓ, τ_c) via L^*, Y^*, α^*), and C^* is given by (2) evaluated at the equilibrium L^*.

For transparency we define two *sufficient-statistics* aggregators at the equilibrium allocation:

$$S(\tau_\ell, \tau_c) \equiv (1-\tau_c)\big(1-\alpha^*\,\tau_\ell\big), \quad \text{(private distortion index)} \tag{5}$$

$$R(\tau_\ell, \tau_c) \equiv \tau_\ell(1-\tau_c)\,\alpha^* + \tau_c, \quad \text{(revenue index per unit of } Y^*) \tag{6}$$

so that $C^* = S(\tau)\, Y^*$ and $G^B = R(\tau)\, Y^*$. In the isoelastic case ($\alpha^* = \beta$ constant), S and R are exactly as in (5)–(6), independent of L^*.

Lemma 1. *Under Assumption 1, for any $(\tau_\ell, \tau_c) \in [0,1)^2$ there exists a unique private equilibrium $L^*(\tau_\ell, \tau_c) > 0$. In the benchmark $\tilde{u}(C, L) = \ln C - \phi(L)$ with ϕ strictly convex and $\lim_{L\to\infty} \phi(L) = +\infty$, L^* is* independent *of (τ_ℓ, τ_c) and solves*

$$\frac{f'(L^*)}{f(L^*)} = \phi'(L^*),$$

so that $Y^* = f(L^*)$ and $k^* = f'(L^*)L^*$ are constants with respect to (τ_ℓ, τ_c).

Standing relations. By Assumption 2, $C^* = S\, Y^*$ and $G^B = k^*\, \widetilde{R}$ with $S = (1-\tau_c)(1-\alpha^*\tau_\ell)$ and $\widetilde{R} = \tau_\ell(1-\tau_c) + \frac{1}{\alpha^*}\tau_c$.

Lemma 2. *Under Assumptions 1–2 and the benchmark $\tilde{u}(C, L) = \ln C - \phi(L)$, $v(G) = G$, the welfare can be written as*

$$W(\tau_\ell, \tau_c; \theta) = \ln(1-\tau_c) + \ln\big(1-\alpha^*\,\tau_\ell\big) \;+\; \theta\, k^* \left[\tau_\ell(1-\tau_c) + \frac{1}{\alpha^*}\tau_c\right] \;-\; \phi(L^*) \;+\; \textit{const.} \tag{7}$$

Hence W is the sum of a strictly concave term in (τ_ℓ, τ_c) (the two logs) and a bilinear term that captures delivered public consumption in expectation.

Lemma 3. *Fix any $s \in (0,1)$ and consider the level set $\mathcal{L}(s) \equiv \{(\tau_\ell, \tau_c) \in (0,1)^2 : \; S(\tau_\ell, \tau_c) = s\}$. Then $\nabla S(\tau_\ell, \tau_c) \neq 0$ on $(0,1)^2$ and, by the implicit function theorem, $\mathcal{L}(s)$ is a C^1 curve. In the isoelastic case ($\alpha^* = \beta$), the slope along $\mathcal{L}(s)$ is*

$$\frac{\mathrm{d}\tau_c}{\mathrm{d}\tau_\ell} = \frac{\alpha^*(1-\tau_c)}{1-\alpha^*\,\tau_\ell} \;>\; 0,$$

so level sets are strictly increasing graphs in (τ_ℓ, τ_c)*-space.*

When $\tilde{u}(C, L) = \ln C - \phi(L)$ and $v(G) = G$ (the benchmark in Section 7), we have

$$W(\tau; \theta) = \ln S(\tau) + \theta\, k^* \, R(\tau) - \phi(L^*) + \text{const}, \tag{8}$$

with $k^* \equiv f'(L^*)L^*$. The first term is concave in (τ_ℓ, τ_c) and the second is linear, so W is concave in rates under standard conditions.

We first record primitives under which the core results admit global (not merely local) statements with closed forms.

Assumption B.
$\tilde{u}(C, L) = \ln C - \phi(L)$ with ϕ convex and C^2; $v(G) = G$; $f : \mathbb{R}_+ \to \mathbb{R}_+$ is C^2, strictly increasing and strictly concave.

Implication.
Under Assumption B, the household's labor choice L^* solves $\max_{L \geq 0}\{\ln f(L) - \phi(L)\}$ and is therefore *independent* of (τ_ℓ, τ_c). Let $Y^* \equiv f(L^*)$, $k^* \equiv f'(L^*)L^*$, and $\alpha^* \equiv k^*/Y^* \in (0, 1)$. Then

$$C^* = (1 - \tau_c)(1 - \alpha^* \, \tau_\ell)\, Y^*, \qquad G^B = \big[\tau_\ell(1 - \tau_c)\alpha^* + \tau_c\big]\, Y^*,$$

and the welfare index (8) reduces to

$$W(\tau; \theta) = \ln\big[(1 - \tau_c)(1 - \alpha^* \tau_\ell)\big] \; + \; \theta\, k^* \, \big[\tau_\ell(1 - \tau_c)\alpha^* + \tau_c\big] \; + \; \text{const}.$$

Lemma 4. *Under Assumption B,* $W(\tau; \theta)$ *is* C^2 *and strictly concave on* $(0, 1)^2$ *in* (τ_ℓ, τ_c)*.*

Proof sketch. $\ln[(1 - \tau_c)(1 - \alpha^* \tau_\ell)] = \ln(1 - \tau_c) + \ln(1 - \alpha^* \tau_\ell)$ has negative diagonal second derivatives and zero cross-partial; the revenue term is linear. Hence the Hessian is diagonal with strictly negative entries on $(0, 1)^2$. □

5 Main Results

We now solve the trust-adjusted Ramsey problem and organize the results in four steps. First, we establish a sharp threshold in trust below which any positive tax reduces welfare (Proposition 1).

Second, once trust clears that threshold, the planner uniquely fixes the scale of taxation while remaining indifferent across a continuum of tax mixes that deliver the same allocation and value—the equivalence frontier (Theorem 1).

Third, we show how tiny, instrument-specific frictions resolve this indifference and select a unique instrument (Proposition 5).

Finally, we summarize the optimality conditions in a trust-adjusted Ramsey rule that is easy to take to data or numerical experiments (Section 5.6). Along the way, we characterize boundary cases, comparative statics in trust, and global concavity to rule out spurious optima.

5.1 A Trust Threshold for Positive Taxation

The basic question is whether it ever makes sense to raise a distortionary tax when citizens doubt delivery. At very low trust, the marginal deadweight loss of taxation dominates the expected benefit of public goods, so the best policy is to refrain from taxing at all. The next result formalizes this intuition by showing that the planner's objective has a clean cutoff in trust: below it, zero taxation is uniquely optimal; above it, some positive revenue is strictly welfare-improving (Proposition 1).

Proposition 1. *Under Assumptions 1–2 and the benchmark* $\tilde{u}(C, L) = \ln C - \phi(L)$*,* $v(G) = G$*, define* $Y^* = f(L^*)$ *at the private equilibrium* L^* *of Lemma 1. Then:*

(i) If $\theta \leq \bar{\theta} \equiv 1/Y^*$*, the unique optimum is zero taxation:* $(\tau_\ell^*, \tau_c^*) = (0, 0)$*.*

(ii) If $\theta > 1/Y^*$*, any welfare-improving policy must raise positive revenue (i.e., some* $\tau_i > 0$ *is strictly beneficial relative to* $(0, 0)$*).*

5.2 Equivalence Frontier: Unique Scale, Indifferent Mix

Once trust is high enough to justify raising revenue, the planner no longer cares which broad instrument does the work—as long as the private distortion is held fixed. This delivers a one-dimensional equivalence frontier in the space of tax rates:[8] many mixes implement the same allocation and value. The frontier pins down the overall scale of taxation, while leaving the mix indeterminate. We characterize this set and show that, along it, private consumption and delivered public consumption move one-for-one with trust in a particularly transparent way (Theorem 1).

Theorem 1. *Maintain Assumptions 1–2 and the benchmark* $\tilde{u} = \ln C - \phi(L)$*,* $v(G) = G$*, and suppose* α^* *is constant (isoelastic technology). If* $\theta > 1/Y^*$*, any interior optimum*

[8] Geometrically, the frontier inherits the shape of level sets of the private-distortion index. In the isoelastic case these level sets are strictly increasing curves; any point on the curve implements the same allocation and value.

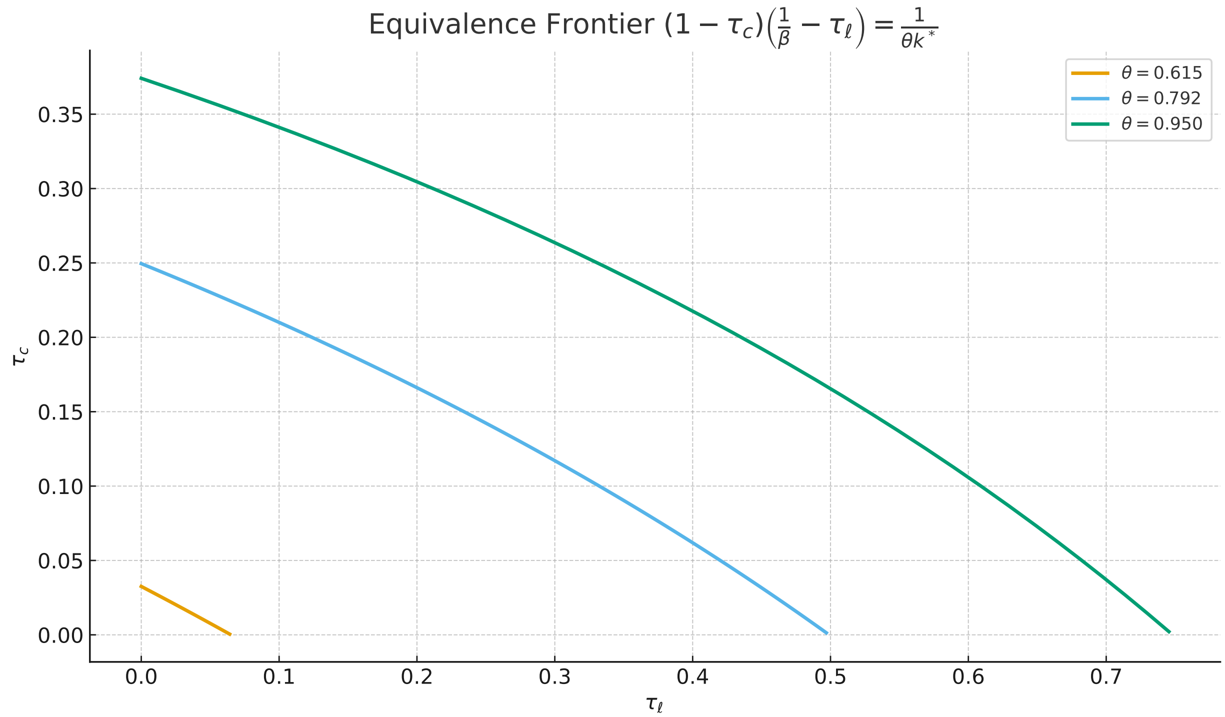

Figure 1: Equivalence frontier in (τ_ℓ, τ_c) space for several trust levels θ (isoelastic, $a = 2$, $\beta = \frac{1}{2}$). Each curve satisfies $(1-\tau_c)\left(\frac{1}{\beta} - \tau_\ell\right) = 1/(\theta k^*)$; all points on a curve implement the same allocation and welfare.

$(\tau_\ell^*, \tau_c^*) \in (0,1)^2$ *satisfies the single equation*

$$(1-\tau_c)\left(\frac{1}{\alpha^*} - \tau_\ell\right) = \frac{1}{\theta k^*}, \tag{9}$$

and the set of interior optima forms a one-dimensional C^1 manifold (the equivalence frontier*) in (τ_ℓ, τ_c)-space. All points on the frontier implement the same allocation and value:*

$$C^*(\theta) = \frac{1}{\theta}, \qquad G^B(\theta) = Y^* - \frac{1}{\theta}, \qquad W^*(\theta) = -\ln\theta + \theta Y^* - \phi(L^*) + const. \tag{10}$$

5.3 KKT characterization, boundaries, and feasibility

Because the feasible policy set includes corners and axis points, a complete description requires the Kuhn–Tucker conditions and a careful look at boundaries. This subsection records the general KKT system, shows when interior solutions exist, and identifies the two polar implementations on the axes. These polar points are convenient for exposition and numerics, and they implement the same allocation and value as any interior point on the frontier.

We work with the box constraints $0 \le \tau_i < 1$, $i \in \{\ell, c\}$.[9]

Let $\mu_i \ge 0$ be the multipliers on $\tau_i \ge 0$ and $\nu_i \ge 0$ on $1 - \tau_i \ge 0$.

Proposition 2. *Under Assumptions 1–2 and the benchmark $\tilde{u} = \ln C - \phi(L)$, $v(G) = G$:*

(i) At any optimum (τ_ℓ^*, τ_c^*),

$$\frac{\partial W}{\partial \tau_\ell}(\tau^*; \theta) + \mu_\ell - \nu_\ell = 0, \qquad \frac{\partial W}{\partial \tau_c}(\tau^*; \theta) + \mu_c - \nu_c = 0.$$

(ii) $\mu_i \, \tau_i^* = 0$ and $\nu_i \, (1 - \tau_i^*) = 0$ for $i \in \{\ell, c\}$.

(iii) $\tau_i^* \in [0, 1)$ and $\mu_i, \nu_i \ge 0$.

(iv) If $(\tau_\ell^*, \tau_c^*) \in (0,1)^2$, then $\mu_i = \nu_i = 0$ and the first-order conditions reduce to

$$-\frac{\alpha^*}{1 - \alpha^* \, \tau_\ell^*} + \theta k^* (1 - \tau_c^*) = 0, \qquad -\frac{1}{1 - \tau_c^*} + \theta k^* \left(\frac{1}{\alpha^*} - \tau_\ell^* \right) = 0,$$

whose solutions form the equivalence frontier (9).

(v) If $\theta > 1/Y^*$, the frontier intersects the open box $(0,1)^2$ and is a nonempty C^1 curve with endpoints on the axes.

Corollary 1. *For $\theta > 1/Y^*$, the two axis points*

$$(\tau_\ell, \tau_c) = \left(0, \; 1 - \frac{1}{\theta Y^*}\right) \quad \textit{and} \quad (\tau_\ell, \tau_c) = \left(\frac{1}{\alpha^*}\left(1 - \frac{1}{\theta Y^*}\right), \; 0\right)$$

satisfy the KKT conditions and implement the same allocation and value (10) *as any interior frontier point.*

5.4 Comparative statics in trust

How do optimal policy objects move with measured trust? Along the frontier, the answer is simple and intuitive: as credibility rises, statutory rates on the selected broad base rise, delivered public consumption increases, and welfare improves at an accelerating pace.[10] The next result gathers these comparative statics and documents their monotonicity and curvature properties in the isoelastic specialization.

[9] We exclude exact $\tau_i = 1$ because the log terms in the objective become ill-defined there and such corner rates are not policy-relevant. Real-world statutory rates are far below unity.

[10] In the isoelastic benchmark with $v(G) = G$ and $\tilde{u}(C, L) = \ln C - \phi(L)$, $W^*(\theta)$ is strictly convex in θ. With concave v, curvature can weaken but monotonicity results remain; see Section 6.

Proposition 3. *In the isoelastic specialization ($\alpha^* = \beta \in (0,1)$) and for $\theta > 1/Y^*$, along the frontier:*

(a) The polar commodity/output rate $\tau_c^*(\theta) = 1 - \frac{1}{\theta Y^*}$ is strictly increasing and concave in θ:

$$\frac{\mathrm{d}\tau_c^*}{\mathrm{d}\theta} = \frac{1}{\theta^2 Y^*} > 0, \qquad \frac{\mathrm{d}^2\tau_c^*}{\mathrm{d}\theta^2} = -\frac{2}{\theta^3 Y^*} < 0.$$

Symmetrically, the polar labor rate $\tau_\ell^*(\theta) = \frac{1}{\alpha^*}\left(1 - \frac{1}{\theta Y^*}\right)$ is strictly increasing and concave.

(b) Delivered public consumption $G^B(\theta) = Y^* - \frac{1}{\theta}$ is strictly increasing and concave in θ:

$$\frac{\mathrm{d}G^B}{\mathrm{d}\theta} = \frac{1}{\theta^2} > 0, \qquad \frac{\mathrm{d}^2 G^B}{\mathrm{d}\theta^2} = -\frac{2}{\theta^3} < 0.$$

(c) Welfare $W^*(\theta) = -\ln\theta + \theta Y^*$ (up to an additive constant) is strictly increasing and strictly convex on $(1/Y^*, \infty)$:

$$\frac{\mathrm{d}W^*}{\mathrm{d}\theta} = -\frac{1}{\theta} + Y^* > 0, \qquad \frac{\mathrm{d}^2 W^*}{\mathrm{d}\theta^2} = \frac{1}{\theta^2} > 0.$$

Proposition 4. *For fixed θ, $W(\tau_\ell, \tau_c; \theta)$ in (7) is the sum of a strictly concave function and a bilinear function, hence concave on $[0,1)^2$. Consequently, any stationary point is a global maximizer. For $\theta \leq 1/Y^*$ the unique maximizer is $(0,0)$; for $\theta > 1/Y^*$ the set of maximizers is precisely the equivalence frontier segment (including its axis endpoints) described by (9).*

5.5 Instrument Selection with Tiny Wedges

Indifference across mixes is knife-edge: in practice, small administrative or salience differences break the tie. We model these as tiny, instrument-specific wedges and show that they uniquely select the revenue workhorse.[11] The conclusion is pragmatic: once trust is sufficient to tax, use the cheaper broad base—unless there is compelling evidence that the alternative base is administratively superior in the relevant range.

Proposition 5. *Let $a_\ell(\tau_\ell)$ and $a_c(\tau_c)$ be C^2, strictly convex with $a_i(0) = a_i'(0) = 0$, representing administrative/salience costs for each instrument. For $\theta > 1/Y^*$, among all points on the frontier (9), the unique optimum minimizes $a_\ell(\tau_\ell) + a_c(\tau_c)$. Under local quadratic*

[11] If administrative or compliance costs are nonconvex (e.g., fixed costs or threshold effects), selection can jump between instruments as the target scale changes. Our quadratic local comparison is conservative and favors the cheaper instrument in a neighborhood of the origin.

costs $a_i(\tau_i) \simeq \frac{\kappa_i}{2}\tau_i^2$ *near the origin, the optimal mix tilts toward the instrument with the smaller curvature* κ_i*; as* $\kappa_j/\kappa_i \to \infty$*, the solution converges to the polar tax using instrument* $i \in \{\ell, c\}$*.*

5.6 A Trust-Adjusted Ramsey Rule

A useful way to summarize optimality is to compare how much private utility is lost at the margin with how much delivered public consumption is gained. The resulting condition looks like the classic Ramsey rule, except that the marginal value of public funds is scaled by trust. This sufficient-statistics form can be read off the data (or a calibration) and helps organize both theory and policy evaluation.

Let $\mathrm{MR}_i \equiv \partial_{\tau_i} G^B$ denote the marginal delivered revenue from instrument $i \in \{\ell, c\}$ and let MEB_i denote the marginal excess burden (the marginal *private* utility loss) from instrument i. Define the *trust-adjusted marginal value of public funds*

$$\mathrm{MVF}(\theta) \;\equiv\; \frac{\theta\, v'(G^B)}{\tilde{u}_C(C, L)}.$$

At any interior optimum with both instruments used,

$$\frac{\mathrm{MEB}_i}{\mathrm{MR}_i} \;=\; \mathrm{MVF}(\theta) \qquad \text{for } i \in \{\ell, c\}, \tag{11}$$

with the usual Kuhn–Tucker inequalities at corners. In the benchmark $\tilde{u}(C, L) = \ln C - \phi(L)$ and $v(G) = G$, one has $\mathrm{MVF}(\theta) = \theta\, C^*$, which equals 1 along the frontier by (10).

6 Extensions and Robustness

The core insights are robust and extend in two natural directions without leaving the one-period benchmark. First, allowing mild non-separability between consumption and leisure preserves the scale result but predictably tilts the preferred mix, providing a simple guide to instrument choice when preferences deviate from additivity. Second, allowing diminishing marginal utility of public goods attenuates the optimal scale while preserving the frontier logic and selection-by-wedges. We present each extension in turn with minimal additional structure.

6.1 Mild non-separability between C and L

Non-separable preferences are empirically relevant and theoretically instructive: when consumption and leisure interact, the same private distortion index can be achieved in subtly different ways by leaning on one instrument or the other. Locally, this breaks frontier indifference and delivers a clear sign test:[12] complements favor the commodity/output base;[13] substitutes favor the labor base. The result provides a disciplined reason to tilt the mix when small preference interactions matter.

We now allow $u(C, L)$ to be C^2, strictly concave, with cross-partial u_{CL} evaluated at the zero-tax allocation $(C_0, L_0) = (Y^*, L^*)$ possibly nonzero, while keeping $v(G) = G$ and Assumptions 1–2. As in Section 3, $C(L, \tau) = (1 - \tau_c)\big(1 - \alpha(L)\tau_\ell\big) f(L)$ with $\alpha(L) = \frac{f'(L)L}{f(L)}$. The private FOC is $F(L, \tau) \equiv u_C(C, L)\, C_L(C, L, \tau) + u_L(C, L) = 0$.

Proposition 6. *Fix $\theta > 1/Y^*$ and the optimal scale S_θ from Theorem 1. Consider infinitesimal policy moves $(\mathrm{d}\tau_\ell, \mathrm{d}\tau_c)$ that* preserve the private-distortion index S to first order*, i.e.*

$$\mathrm{d}S = S_{\tau_\ell}\, \mathrm{d}\tau_\ell + S_{\tau_c}\, \mathrm{d}\tau_c = 0 \quad \Longleftrightarrow \quad \mathrm{d}\tau_c = -\alpha^*\, \mathrm{d}\tau_\ell \quad \text{at } (0, 0).$$

Then the second-order welfare change at the zero-tax allocation satisfies

$$\mathrm{d}^2 W \;=\; -\,\Xi\, u_{CL}(C_0, L_0)\, \big(\mathrm{d}\tau_\ell\big)^2 \;+\; o\big(\|\mathrm{d}\tau\|^2\big)\,, \tag{12}$$

where $\Xi > 0$ depends only on primitives (curvature of u and f) at (C_0, L_0).[14] *In particular:*

- *If $u_{CL}(C_0, L_0) < 0$ (consumption and leisure are complements), $\mathrm{d}^2 W < 0$ when increasing τ_ℓ and compensating τ_c to keep S fixed, so the planner prefers tilting the mix toward the commodity/output tax.*
- *If $u_{CL}(C_0, L_0) > 0$ (substitutes), the planner prefers tilting toward the labor-income tax.*

Thus, index-based indifference in the separable benchmark is robust in scale *but breaks predictably in* mix *under mild non-separability.*

[12]For GHH-type preferences (additive in $C - \psi(L)$), the cross-partial u_{CL} is zero at the benchmark, so the local tilt disappears. Small deviations from GHH restore a definite tilt with the sign given by u_{CL}.

[13]We abstract from capital, savings, and intertemporal choice to isolate the trust channel.

[14]One convenient representation is $\Xi = \frac{\left(C_L(C_0, L_0)\right)^2}{-F_L(C_0, 0)}$, with $C_L = f'(L_0) > 0$ and $F_L < 0$ by strict concavity; see Appendix.

6.2 Concave public-good utility

Public goods often exhibit diminishing marginal utility once a minimum service level is achieved. Incorporating this curvature does not undermine the logic of the benchmark. It shifts the trust threshold in a transparent way and reduces the optimal scale[15] for any given trust while leaving the equivalence frontier and the selection-by-wedges conclusion intact. We provide the characterization and the comparative statics.

Let $v : \mathbb{R}_+ \to \mathbb{R}$ be C^2, nondecreasing and concave ($v' > 0$, $v'' \leq 0$), while the private part remains $\tilde{u}(C, L) = \ln C - \phi(L)$ as in the benchmark, and Assumptions 1–2 hold.

Proposition 7. *At the origin, the directional derivatives satisfy*

$$\left.\frac{\partial W}{\partial \tau_\ell}\right|_{(0,0)} = -\alpha^* + \theta\, v'(0)\, k^*, \qquad \left.\frac{\partial W}{\partial \tau_c}\right|_{(0,0)} = -1 + \theta\, v'(0)\, Y^*.$$

Hence the trust threshold generalizes to

$$\bar{\theta}_v \;=\; \frac{1}{v'(0)\, Y^*}.$$

If $\theta \leq \bar{\theta}_v$, $(0,0)$ is optimal; if $\theta > \bar{\theta}_v$, some positive tax is strictly welfare-improving.

Theorem 2. *In the isoelastic case ($\alpha^* = \beta$ constant), W depends on (τ_ℓ, τ_c) only through $S(\tau) = (1-\tau_c)(1-\alpha^*\tau_\ell)$ and $G^B(\tau) = (1-S(\tau))Y^*$. For $\theta > \bar{\theta}_v$, the optimal* scale *$S_\theta \in (0,1)$ is characterized uniquely by*

$$\frac{1}{S_\theta} \;=\; \theta\, v'\big((1-S_\theta)Y^*\big)\, Y^*, \tag{13}$$

and the mix *is indeterminate: any (τ_ℓ, τ_c) satisfying $S(\tau) = S_\theta$ is optimal.*

Proposition 8. *Let S_θ solve* (13). *Then*

$$\frac{\mathrm{d}S_\theta}{\mathrm{d}\theta} \;=\; \frac{v'\big((1-S_\theta)Y^*\big)\, Y^*}{-\dfrac{1}{S_\theta^2} \;+\; \theta\, v''\big((1-S_\theta)Y^*\big)\, (Y^*)^2} \;<\; 0,$$

because $v' > 0$ and $v'' \leq 0$ imply the denominator is strictly negative. Therefore $C^*(\theta) = S_\theta Y^*$ is strictly decreasing and $G^B(\theta) = (1-S_\theta)Y^*$ strictly increasing in θ.

Remark 2. Concavity of v attenuates the optimal tax scale relative to the linear case $v(G) = G$, but the equivalence-frontier logic (mix indifference at a fixed S_θ) and wedge-based instrument selection remain unchanged.

[15] Calibrating $v'(0)$ amounts to choosing the social value of the first unit of public services (e.g., basic security or primary health). Higher $v'(0)$ lowers the threshold and justifies taxation at lower measured trust.

7 Isoelastic Specialization

To make the logic fully transparent and to facilitate figures, we specialize to a simple isoelastic technology[16] and a log-linear private utility. This yields closed-form expressions for the trust threshold, polar tax schedules, delivered public consumption, and welfare. The specialization is not needed for the theory but is helpful for calibration, teaching, and quick policy analytics; the figures visualize how optimal objects rise with trust and where the frontier lies in rate space.

Let $f(L) = aL^\beta$ with $a > 0$ and $\beta \in (0, 1)$; take $\tilde{u}(C, L) = \ln C - \frac{1}{2}L^2$ and $v(G) = G$. Then $\alpha^* = \beta$, the equilibrium input is $L^* = \sqrt{\beta}$, output $Y^* = a\,\beta^{\beta/2}$, and $k^* = \beta Y^*$. The sufficient-statistics reduce to

$$S(\tau) = (1 - \tau_c)(1 - \beta\tau_\ell), \qquad R(\tau) = \beta\tau_\ell(1 - \tau_c) + \tau_c. \tag{14}$$

Threshold. The trust threshold equals

$$\bar{\theta} = \frac{1}{Y^*} = \frac{1}{a\,\beta^{\beta/2}}. \tag{15}$$

Frontier and polar implementations. For $\theta > \bar{\theta}$, the frontier is

$$(1 - \tau_c)\Big(\tfrac{1}{\beta} - \tau_\ell\Big) = \frac{1}{\theta k^*}. \tag{16}$$

Two polar optima implement the same allocation and welfare:

$$\tau_c^*(\theta) = 1 - \frac{1}{\theta Y^*}, \qquad \tau_\ell^*(\theta) = \frac{1}{\beta}\left(1 - \frac{1}{\theta Y^*}\right). \tag{17}$$

Along the frontier $C^*(\theta) = 1/\theta$, $G^B(\theta) = Y^* - \frac{1}{\theta}$, and $W^*(\theta) = -\ln\theta + \theta Y^* - \frac{1}{2}(L^*)^2 + \text{const}$.

8 Policy Implications

The analysis yields a compact policy map for environments where the conversion of revenue into public services is uncertain.

When measured trust is below the threshold, any distortionary tax worsens welfare. The right lever is not the rate but credibility: actions that make delivery more visible, verifiable, and timely. Examples include transparent procurement and publication of delivery milestones;

[16] The figures use $a = 2$ and $\beta = \frac{1}{2}$ purely for illustration. Any monotone rescaling of a or moderate change in β preserves the shapes and the qualitative comparative statics.

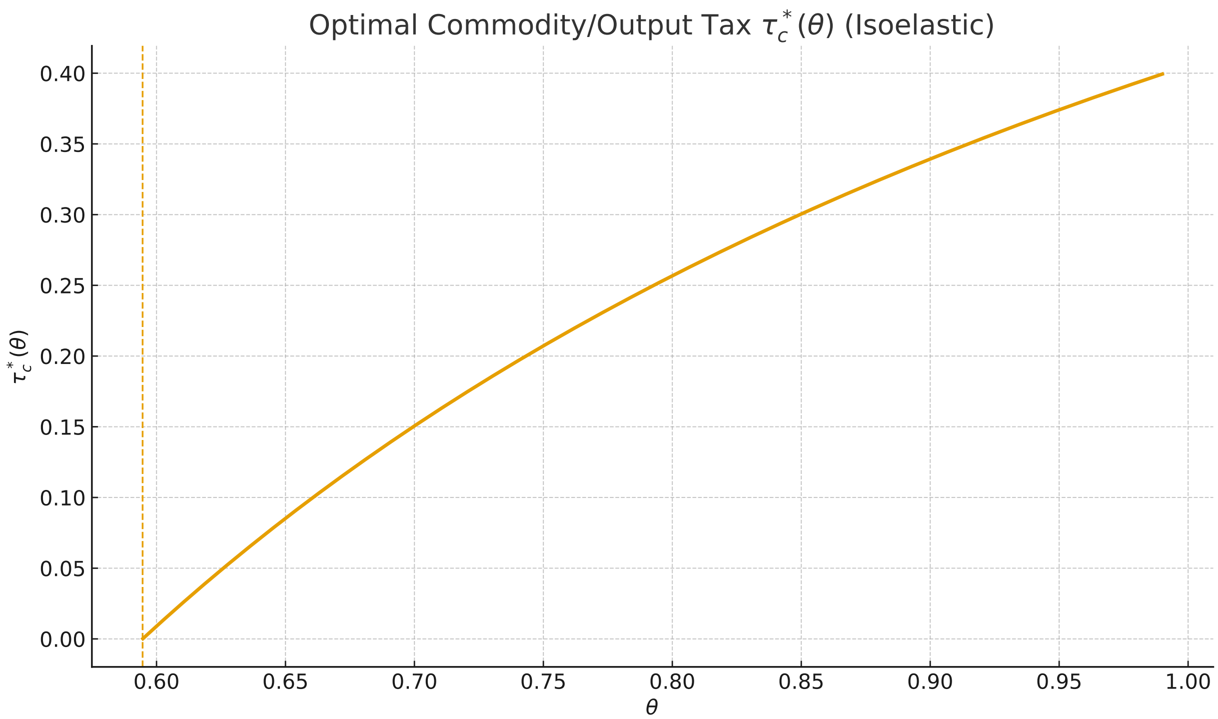

Figure 2: Optimal commodity/output tax $\tau_c^*(\theta)$ in the isoelastic specialization ($a = 2$, $\beta = \frac{1}{2}$). The dashed line marks the trust threshold $\bar{\theta} = 1/Y^*$.

third-party audits; narrow, highly visible pilot projects; and clear ex ante commitments about scope and timelines. In our benchmark, such actions raise the effective trust parameter and move the economy into the region where positive taxation is justified.

Above the threshold, the planner fixes the scale and is indifferent over the mix in the frictionless benchmark. In practice, small administrative and compliance differences matter: broad bases collected at source (such as a well-designed VAT or uniform sales tax) typically dominate on collection cost and enforceability. The model provides the normative justification: when tiny wedges break indifference, use the instrument with the lower wedge. Where the labor base is demonstrably cheaper to collect (e.g., highly formalized payroll systems with strong withholding), tilt accordingly.

As credibility improves, optimal statutory rates on the selected base rise monotonically, delivered public consumption increases one-for-one with the contraction in private consumption, and welfare improves at an accelerating pace. This suggests a practical rule: calibrate or proxy trust (from surveys, delivery audits, or outcome-tracking dashboards) and map it into a target tax-to-GDP for the selected base. The isoelastic specialization provides ready-to-use formulas and visuals for communication and planning.

To lift trust efficiently, choose projects and delivery mechanisms with short feedback

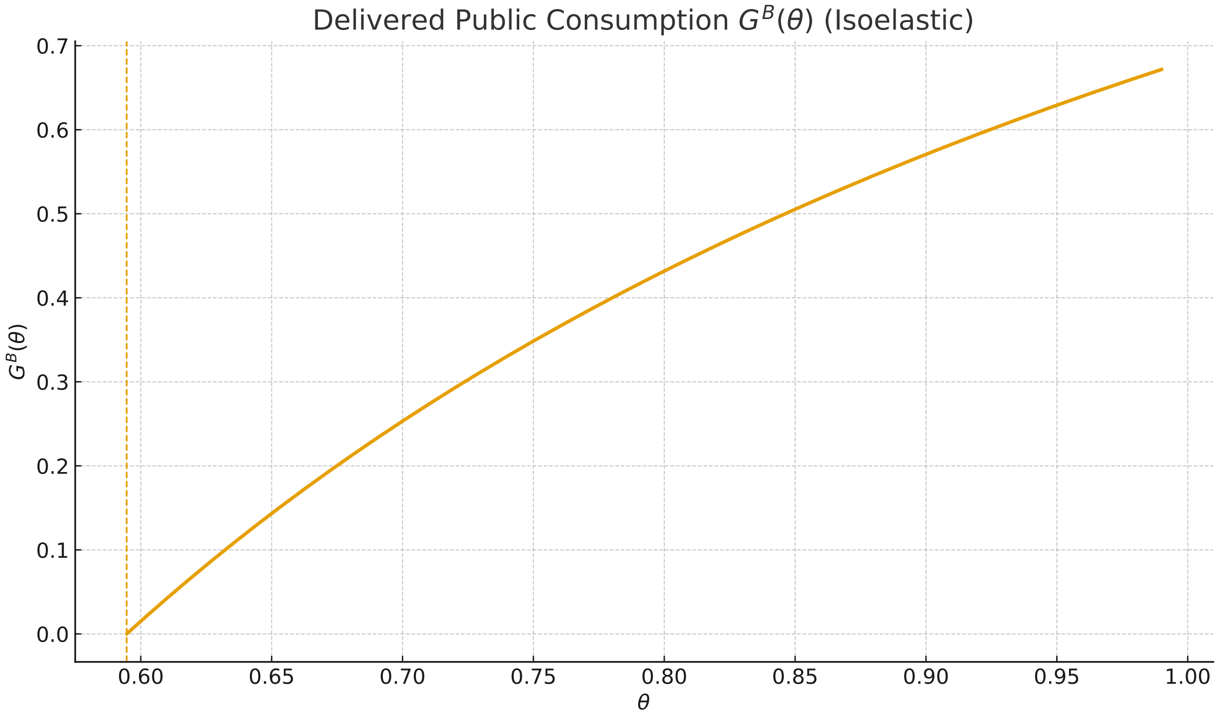

Figure 3: Delivered public consumption $G^B(\theta) = Y^* - 1/\theta$ along the equivalence frontier (isoelastic case). The dashed line marks $\bar{\theta} = 1/Y^*$.

loops and high observability (e.g., maintenance with visible outputs, digital public services with user-side logs). Where feasible, publish contract-level data and simple indicators of completion. Even in a static benchmark, these choices act like raising the effective conversion rate for the marginal unit of revenue.

When policymakers face a choice between two bases with similar statutory reach, a small diagnostic calculation—how much administrative/salience wedge is needed to prefer one base over the other at the target scale—can be decisive. If that wedge is plausibly positive (e.g., difference in enforcement costs or evasion exposure), the model's selection result gives clear guidance.

For public buy-in, the policy message can be made simple: "We raise rates only when we can credibly deliver. As trust increases, we scale gradually and use the cheapest base to collect." This aligns incentives on both sides—government commits to visibility and delivery; citizens observe and update trust.

The benchmark is static and representative-agent. It abstracts from distributional objectives, sectoral specifics, and explicit compliance responses. These omissions are deliberate and help isolate the trust channel. Extensions can add distributional weights, heterogeneity in trust, or simple compliance responses without compromising the central policy messages.

Figure 4: Welfare $W^*(\theta) = -\ln\theta + \theta Y^* - \frac{1}{2}(L^*)^2$ (plotted up to an additive constant) along the frontier. The dashed line marks $\bar{\theta} = 1/Y^*$.

9 Conclusion

This paper places public trust at the center of the optimal tax problem.[17] In a static Ramsey benchmark with two broad instruments, trust acts as a primitive that determines whether to tax at all, how much to tax when it is worthwhile, and which instrument should carry the load when tiny practical wedges matter. The theory yields a threshold for taxing, an equivalence frontier for the mix, and a trust-adjusted Ramsey rule in sufficient statistics. The resulting policy map is simple: build credibility first; once trust is sufficient, keep the base broad and let measured trust determine the scale.

Acknowledgments

We thank Jacob Gershman, Arthur Izgarshev, David Li, and Anastasia Makhmudova for helpful comments. All remaining errors are our own.

[17]External validity will depend on heterogeneity (in incomes and trust), political frictions, and sectoral structure. These extensions are compatible with the benchmark but left for future work.

Data Availability

No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Funding

This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

A Proofs

This appendix provides formal proofs and omitted derivations.

Proof of Lemma 1

Fix (τ_ℓ, τ_c). The household maximizes $L \mapsto \tilde{u}\big(C(L), L\big)$ where, by firm optimality and the budget, $C(L) = (1-\tau_c)\big(1-\alpha(L)\tau_\ell\big)f(L)$ with $\alpha(L) \equiv f'(L)L/f(L) \in (0,1)$. Under Assumption 1(ii), $\tilde{u}$ is strictly concave; under (i), f is concave and increasing; hence $L \mapsto \tilde{u}(C(L), L)$ is strictly concave provided u_1 is concave and nondecreasing in C (composition with a concave map) and u_2 is strictly concave in L. Coercivity follows from $\lim_{L\to\infty} \tilde{u}(C(L), L) = -\infty$ (e.g. when $\tilde{u} = \ln C - \phi(L)$ with ϕ strictly convex and unbounded), and the objective tends to $-\infty$ as $L \downarrow 0$ whenever $f(0) = 0$ and $\tilde{u}(C, \cdot)$ penalizes vanishing C (e.g. $\ln C \to -\infty$), implying an interior maximizer. Strict concavity delivers uniqueness.

For the benchmark $\tilde{u}(C, L) = \ln C - \phi(L)$, we have

$$\max_{L>0} \ln\big(S(\tau_\ell, \tau_c)\, f(L)\big) - \phi(L) = \ln S(\tau_\ell, \tau_c) + \max_{L>0} \big\{ \ln f(L) - \phi(L) \big\},$$

so the argmax L^* solves $\big(f'(L)/f(L)\big) - \phi'(L) = 0$ and is independent of (τ_ℓ, τ_c). Setting $Y^* = f(L^*)$ and $k^* = f'(L^*)L^*$ yields the stated constants. □

Proof of Lemma 2

Under Assumption 2, $C^* = S\,Y^*$ and $G^B = (1-S)Y^*$ with $S = S(\tau_\ell, \tau_c)$, Y^* constant. With $\tilde{u}(C, L) = \ln C - \phi(L)$ and $v(G) = G$,

$$W(\tau_\ell, \tau_c; \theta) = \ln(SY^*) - \phi(L^*) + \theta\,(1-S)Y^* + \text{const} = \ln S + \theta k^*(1-S) - \phi(L^*) + \text{const},$$

since $k^* = f'(L^*)L^*$ and $Y^* = \frac{k^*}{\alpha^*}$. Hence W depends on (τ_ℓ, τ_c) only through $S(\tau_\ell, \tau_c)$. Finally,

$$\ln S(\tau_\ell, \tau_c) = \ln(1 - \tau_c) + \ln\bigl(1 - \alpha^* \tau_\ell\bigr),$$

a sum of concave functions on $[0, 1)^2$ (each term is the log of a positive affine function), hence strictly concave. □

Proof of Lemma 3

Compute the gradient

$$\nabla S(\tau_\ell, \tau_c) = \bigl(- \alpha^*(1 - \tau_c),\ -(1 - \alpha^* \tau_\ell)\bigr),$$

which is nonzero on $(0, 1)^2$. By the implicit function theorem, the level set $\mathcal{L}(s) = \{S = s\}$ is a C^1 curve. In the isoelastic case $\alpha^* = \beta$, differentiating $S(\tau_\ell, \tau_c) = (1 - \tau_c)(1 - \alpha^* \tau_\ell) = s$ yields

$$-(1 - \alpha^* \tau_\ell)\, \mathrm{d}\tau_c - \alpha^*(1 - \tau_c)\, \mathrm{d}\tau_\ell = 0 \ \ \Rightarrow\ \ \frac{\mathrm{d}\tau_c}{\mathrm{d}\tau_\ell} = \frac{\alpha^*(1 - \tau_c)}{1 - \alpha^* \tau_\ell} > 0.$$

Thus $\mathcal{L}(s)$ is strictly increasing in (τ_ℓ, τ_c)-space. □

Preliminaries under Assumption B. Write $S(\tau) \equiv (1 - \tau_c)(1 - \alpha^*\tau_\ell)$ and $R(\tau) \equiv \tau_\ell(1-\tau_c)\alpha^* + \tau_c$. By Assumption B, L^* (hence Y^*, k^*, α^*) is independent of (τ_ℓ, τ_c). Therefore

$$W(\tau; \theta) = \ln S(\tau) + \theta k^* R(\tau) + \text{const},$$

which is concave by Lemma 4.

Proof of Proposition 1

From (7), the partial derivatives at the origin are

$$\left.\frac{\partial W}{\partial \tau_\ell}\right|_{(0,0)} = -\alpha^* + \theta k^*, \qquad \left.\frac{\partial W}{\partial \tau_c}\right|_{(0,0)} = -1 + \theta \frac{k^*}{\alpha^*} = -1 + \theta Y^*.$$

Hence, if $\theta \le 1/Y^*$, both one-sided partials are ≤ 0, so no feasible *revenue-raising* movement from $(0, 0)$ can increase W. Since W is continuous on the compact $[0, 1]^2$ and (by the benchmark) strictly decreases in each log term initially, $(0, 0)$ is a global maximizer. If $\theta > 1/Y^*$, then $\partial_{\tau_c} W|_{(0,0)} > 0$, so increasing τ_c slightly (raising revenue) increases W; in particular, any improvement requires strictly positive revenue. □

Proof of Theorem 1

Compute the first-order conditions from (7):

$$\frac{\partial W}{\partial \tau_\ell} = -\frac{\alpha^*}{1-\alpha^*\,\tau_\ell} + \theta k^*(1-\tau_c) \;=\; 0,$$
$$\frac{\partial W}{\partial \tau_c} = -\frac{1}{1-\tau_c} + \theta k^*\left(\frac{1}{\alpha^*} - \tau_\ell\right) \;=\; 0.$$

Rearranging the two equations and multiplying sides yields

$$(1-\tau_c)\left(\frac{1}{\alpha^*} - \tau_\ell\right) = \frac{1}{\theta k^*},$$

which is (9). For $\theta > 1/Y^*$, feasibility of (9) with $(\tau_\ell, \tau_c) \in (0,1)^2$ is immediate (e.g., set $\tau_\ell = 0$ and solve for τ_c). The implicit function theorem (as in Lemma 3) implies the solution set is a C^1 curve.

To compute the implemented allocation along the frontier, note that

$$S(\tau) = (1-\tau_c)(1-\alpha^*\tau_\ell) = \alpha^*\,(1-\tau_c)\left(\frac{1}{\alpha^*} - \tau_\ell\right) = \frac{\alpha^*}{\theta k^*} = \frac{1}{\theta Y^*}.$$

Hence $C^*(\theta) = S\,Y^* = 1/\theta$ and $G^B(\theta) = Y^* - C^*(\theta) = Y^* - 1/\theta$. Substituting into (7) at any frontier point yields

$$W^*(\theta) = \ln C^*(\theta) - \phi(L^*) + \theta\, G^B(\theta) + \text{const} = -\ln\theta + \theta Y^* - \phi(L^*) + \text{const},$$

which is independent of the mix (τ_ℓ, τ_c) on (9). □

Proof of Proposition 2

Parts (i)–(iii) are the standard KKT conditions for the box-constrained maximization of the concave W. For (iv), interior optimality implies $\mu_i = \nu_i = 0$ and stationarity reduces to the FOCs computed from (7), yielding the frontier equations. For (v), if $\theta > 1/Y^*$ then $1/(\theta k^*) \in (0, 1/\alpha^*)$ since $k^* = \alpha^* Y^*$; the curve $(1-\tau_c)(1/\alpha^* - \tau_\ell) = 1/(\theta k^*)$ intersects each axis at the finite points reported in Corollary 1, so it enters $(0,1)^2$. By Lemma 3 and $\nabla S \neq 0$ on $(0,1)^2$, the frontier is a nonempty C^1 curve with those endpoints. □

Proof of Corollary 1

Evaluate the FOCs at $(\tau_\ell, \tau_c) = \big(0,\, 1 - \frac{1}{\theta Y^*}\big)$: $\partial_{\tau_c} W = -\frac{1}{1-\tau_c} + \theta k^*(1/\alpha^*) = -\theta Y^* + \theta Y^* = 0$ and $\partial_{\tau_\ell} W = -\alpha^* + \theta k^*(1-\tau_c) = -\alpha^* + \theta k^*(1/(\theta Y^*)) = 0$. With $\nu_i = 0$ (both rates strictly below 1), stationarity holds with $\mu_i = 0$ and complementary slackness is satisfied since $\tau_\ell = 0$. The other polar point is analogous. Both lie on (9), hence implement (10). □

Proof of Proposition 3

(a) and (b) are direct differentiations of $\tau_c^*(\theta) = 1 - 1/(\theta Y^*)$, $\tau_\ell^*(\theta) = (1/\alpha^*)(1 - 1/(\theta Y^*))$, and $G^B(\theta) = Y^* - 1/\theta$. (c) Differentiating $W^*(\theta) = -\ln\theta + \theta Y^*$ yields $W'^*(\theta) = -1/\theta + Y^*$ and $W''^*(\theta) = 1/\theta^2 > 0$; since $\theta > 1/Y^*$, $W'^*(\theta) > 0$. □

Proof of Proposition 4

From (7), $W(\tau_\ell, \tau_c; \theta) = \ln(1-\tau_c) + \ln(1-\alpha^*\tau_\ell) + \theta k^*\big[\tau_\ell(1-\tau_c) + (1/\alpha^*)\tau_c\big] - \phi(L^*) + \text{const.}$ The sum of the two log terms is strictly concave on $[0,1)^2$; the last bracket is bilinear, hence affine in each coordinate and thus preserves concavity when added. Therefore W is concave on $[0,1)^2$ and any stationary point is globally optimal. If $\theta \le 1/Y^*$, Proposition 1 gives $(0,0)$ as the unique maximizer. If $\theta > 1/Y^*$, the interior stationarity equations define the C^1 frontier (9); including the axis endpoints (Corollary 1) yields the full set of maximizers. □

Proof of Proposition 5

For $\theta > 1/Y^*$, the set of interior optima without costs is the frontier $\mathcal{F}(\theta) \equiv \{(\tau_\ell, \tau_c) : (1-\tau_c)(\frac{1}{\alpha^*} - \tau_\ell) = 1/(\theta k^*)\}$. Adding instrument-specific costs $a_\ell(\tau_\ell) + a_c(\tau_c)$ changes W by a strictly convex penalty along $\mathcal{F}(\theta)$. Therefore the maximizer of $W - a_\ell - a_c$ on $\mathcal{F}(\theta)$ is equivalent to

$$\min_{(\tau_\ell,\tau_c)\in\mathcal{F}(\theta)} \; a_\ell(\tau_\ell) + a_c(\tau_c),$$

which has a unique solution by strict convexity and the fact that $\mathcal{F}(\theta)$ is a C^1 curve. The quadratic local-comparison claim follows by linearizing the frontier near a point with small rates (so that $(1-\tau_c) \approx 1$ and $(1/\alpha^* - \tau_\ell) \approx 1/\alpha^*$) and solving the quadratic program $\min \frac{\kappa_\ell}{2}\tau_\ell^2 + \frac{\kappa_c}{2}\tau_c^2$ subject to $\tau_\ell + \frac{1}{\alpha^*}\tau_c = \text{const}$: the solution tilts toward the instrument with smaller curvature; the polar limit obtains as the curvature ratio diverges. □

Proof of (11) (trust-adjusted Ramsey rule)

Write the planner's objective as $W = \tilde{u}(C, L) + \theta v(G^B)$ and consider a small change in τ_i. Let $\mathrm{MR}_i \equiv \partial_{\tau_i} G^B$ and define the marginal private utility loss (excess burden) as $\mathrm{MEB}_i \equiv -(\tilde{u}_C\, \partial_{\tau_i} C + \tilde{u}_L\, \partial_{\tau_i} L)$. At an interior optimum, the total derivative satisfies $0 = \partial_{\tau_i} W = -\mathrm{MEB}_i + \theta v'(G^B)\,\mathrm{MR}_i$. Dividing by $\tilde{u}_C > 0$ gives

$$\frac{\mathrm{MEB}_i}{\mathrm{MR}_i} = \frac{\theta v'(G^B)}{\tilde{u}_C(C, L)} \;=\; \mathrm{MVF}(\theta),$$

establishing (11). In the benchmark $\tilde{u}(C, L) = \ln C - \phi(L)$ and $v(G) = G$, we have $\tilde{u}_C = 1/C$ and $v'(G) = 1$, so $\mathrm{MVF}(\theta) = \theta C$; along the frontier $C^*(\theta) = 1/\theta$ by (10), hence $\mathrm{MVF}(\theta) = 1$. □

Proof of Proposition 6

Work at the zero-tax allocation $(C_0, L_0) = (Y^*, L^*)$. Let $F(L, \tau) \equiv u_C(C, L)\, C_L(C, L, \tau) + u_L(C, L)$ denote the private FOC, with $F(C_0, L_0, 0) = 0$ and $F_L < 0$ by strict concavity. By the implicit function theorem, $L(\tau)$ is C^1 near $\tau = 0$ and

$$\mathrm{d}L \;=\; -\,\frac{F_{\tau_\ell}\,\mathrm{d}\tau_\ell + F_{\tau_c}\,\mathrm{d}\tau_c}{F_L}\,.$$

Along directions that preserve S to first order at the origin, $\mathrm{d}\tau_c = -\alpha^*\,\mathrm{d}\tau_\ell$. Moreover $C_{\tau_\ell}(C_0, L_0, 0) = -f'(L_0)L_0 = -k^*$ and $C_{\tau_c}(C_0, L_0, 0) = -f(L_0) = -Y^*$, whence the *direct* change in C cancels: $C_{\tau_\ell}\mathrm{d}\tau_\ell + C_{\tau_c}\mathrm{d}\tau_c = -k^*\,\mathrm{d}\tau_\ell + \alpha^* Y^*\,\mathrm{d}\tau_\ell = 0$ (since $k^* = \alpha^* Y^*$). Similarly, the public-good term changes only at $O(\|\mathrm{d}\tau\|^2)$: G^B is linear in τ times bases, so its variation via L is second order at the origin.

Hence the first-order welfare change vanishes and the leading term is second order via the induced $\mathrm{d}L$:

$$\mathrm{d}W \;=\; u_C\, C_L\,\mathrm{d}L + u_L\,\mathrm{d}L + o(\|\mathrm{d}\tau\|) \;=\; \big(u_C C_L + u_L\big)\mathrm{d}L + o(\|\mathrm{d}\tau\|) = o(\|\mathrm{d}\tau\|),$$

using $F = 0$ at $(C_0, L_0, 0)$. A standard second-order expansion (or differentiating the envelope condition) yields

$$\mathrm{d}^2 W \;=\; -\,\frac{\big(u_{CL}(C_0, L_0)\big)\,\big(C_L(C_0, L_0)\big)^2}{-F_L(C_0, L_0, 0)}\,\big(\mathrm{d}\tau_\ell\big)^2 \;+\; o\big(\|\mathrm{d}\tau\|^2\big)\,,$$

after substituting $\mathrm{d}\tau_c = -\alpha^*\,\mathrm{d}\tau_\ell$ and collecting terms; see, e.g., the quadratic form obtained from the implicit-function reduction of $L(\tau)$. Set $\Xi \equiv \frac{(C_L)^2}{-F_L} > 0$ to obtain (12) and the sign conclusions. $\square$

Proof of Proposition 7

With v concave, $W(\tau;\theta) = \ln(1-\tau_c) + \ln(1-\alpha^*\tau_\ell) + \theta\, v\big(k^*[\tau_\ell(1-\tau_c)] + \frac{k^*}{\alpha^*}\tau_c\big) - \phi(L^*) + \text{const}$ near the origin. Differentiating at $(0,0)$ gives the stated partials: $\partial_{\tau_\ell} W|_{(0,0)} = -\alpha^* + \theta v'(0)k^*$ and $\partial_{\tau_c} W|_{(0,0)} = -1 + \theta v'(0)Y^*$. Thus the threshold is $\bar{\theta}_v = 1/(v'(0)Y^*)$ by the same concavity argument as in Proposition 1. $\square$

Proof of Theorem 2

Under $\tilde{u} = \ln C - \phi(L)$, $C^* = SY^*$ with $S = (1-\tau_c)(1-\alpha^*\tau_\ell)$, and $G^B = (1-S)Y^*$; hence $W(\tau;\theta) = \ln S + \theta\, v\big((1-S)Y^*\big) + \text{const}$ depends on (τ_ℓ, τ_c) only through S. Maximizing over $S \in (0,1)$ yields the first-order condition $0 = \partial W/\partial S = 1/S - \theta\, v'\big((1-S)Y^*\big)\, Y^*$, which uniquely determines S_θ because the left-hand side is strictly decreasing in S (since $v'' \leq 0$). Any (τ_ℓ, τ_c) with $S(\tau) = S_\theta$ is optimal, proving mix indifference. $\square$

Proof of Proposition 8

Let $F(S,\theta) \equiv \frac{1}{S} - \theta\, v'\big((1-S)Y^*\big)\, Y^*$. At S_θ we have $F(S_\theta, \theta) = 0$ and $\partial_\theta F = -\, v'\big((1-S_\theta)Y^*\big)\, Y^* < 0$, while $\partial_S F = -\frac{1}{S_\theta^2} + \theta\, v''\big((1-S_\theta)Y^*\big)\, (Y^*)^2 < 0$ because $v'' \leq 0$. By the implicit function theorem, $\mathrm{d}S_\theta/\mathrm{d}\theta = -(\partial_\theta F)/(\partial_S F) < 0$. Therefore $C^*(\theta) = S_\theta Y^*$ decreases and $G^B(\theta) = (1-S_\theta)Y^*$ increases with θ. $\square$

References

Atkinson, A. B. and Stiglitz, J. E. (1976). The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics*, 6(1-2):55–75.

Barro, R. J. (1986). Reputation in a model of monetary policy with incomplete information. *Journal of Monetary Economics*, 17(1):3–20.

Barro, R. J. and Gordon, D. B. (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1):101–121.

Chari, V. V. and Kehoe, P. J. (1990). Sustainable plans. *Journal of Political Economy*, 98(4):783–802.

Debortoli, D. and Nunes, R. (2010). Fiscal policy under loose commitment. *Journal of Economic Theory*, 145(3):1005–1032.

Diamond, P. A. and Mirrlees, J. A. (1971a). Optimal taxation and public production i: Production efficiency. *American Economic Review*, 61(1):8–27.

Diamond, P. A. and Mirrlees, J. A. (1971b). Optimal taxation and public production ii: Tax rules. *American Economic Review*, 61(3):261–278.

Ebrill, L., Keen, M., Bodin, J.-P., and Summers, V. (2001). *The Modern VAT*. International Monetary Fund, Washington, DC.

Kydland, F. E. and Prescott, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491.

Lu, Y. K. (2013). Optimal policy with endogenous credibility. *Journal of Economic Theory*, 148(2):426–455.

Phelan, C. (2006). Public trust and government betrayal. *Journal of Economic Theory*, 130(1):27–65.

Pomeranz, D. (2015). No taxation without information: Deterrence and self-enforcement in the VAT. *American Economic Review*, 105(8):2539–2569.

Ramsey, F. P. (1927). A contribution to the theory of taxation. *The Economic Journal*, 37(145):47–61.

Scholz, J. T. and Lubell, M. (1998). Trust and taxpaying: Testing the heuristic approach to collective action. *American Journal of Political Science*, 42(2):398–417.

Slemrod, J. and Yitzhaki, S. (2002). Tax avoidance, evasion, and administration. In Auerbach, A. J. and Feldstein, M., editors, *Handbook of Public Economics*, volume 3, pages 1423–1470. Elsevier.

Torgler, B. (2007). *Tax Compliance and Tax Morale: A Theoretical and Empirical Analysis*. Edward Elgar, Cheltenham.