Public Good Provision and Optimal Taxation in a Hidden Savings World

Alberto Ramirez de Aguilar

Universidad Anáhuac

ITAM Alumni Conference, August 2026

Motivation

- In economies with widespread financial informality, the government cannot observe individual income.
- It can, however, observe formal financial activity: **deposits in the banking system**.
- This creates an incentive to hide income by saving in cash ("the couch") instead of the bank.

Motivation

- In economies with widespread financial informality, the government cannot observe individual income.
- It can, however, observe formal financial activity: **deposits in the banking system**.
- This creates an incentive to hide income by saving in cash ("the couch") instead of the bank.
- **Question:** how should a government tax savings, and how generous should public good provision be, when taxation itself can push agents into financial informality?

This Paper

- Two-period economy: agents earn income when young, save for when old, and consume savings plus a publicly provided good.
- Agents can save in a **bank** (taxed, return $R > 1$) or in the **couch** (untaxed, return 1).
- The government can only tax *observed bank savings*, not income directly.
- **Main result:** the welfare-maximizing tax is **concave and flat beyond a threshold** — it deters evasion at the top and tolerates financial informality at the bottom.
- This is the *opposite* of the benchmark (full-information) prescription, which is progressive.
- **Roadmap:** (1) Model → (2) Benchmark (full information) → (3) Main model (hidden savings) → (4) Quantitative illustrations & policy takeaways.

Model

- Continuum of agents with income type $e \in E = \{e_1, \ldots, e_n\}$, $\Pr(e_j) = \pi_j$.
- **Young:** receive e, pay taxes, choose bank savings $b \geq 0$ and couch savings $b^c \geq 0$.
- **Old:** consume savings (with interest) plus public good g.
- An agent of type e solves:

$$\max_{b,\, b^c \geq 0} \quad u\Big(\underbrace{e - b - b^c - \tau(\cdot)}_{c^{young}}\Big) + \beta\, u\Big(\underbrace{Rb + b^c + \psi g}_{c^{old}}\Big),$$

- Public good provision is funded by tax collection: $g = \sum_j \tau(\cdot)\, \pi_j$.
- Government **commits** to τ first, to maximize utilitarian welfare. Assume $\beta R = 1$, $R \leq \psi$.

Benchmark: The Government Observes Income Details

- Suppose (only for this slide) the government observes e directly and taxes $\tau(e)$.
- Since $R > 1$, the couch is strictly dominated: **nobody hides income**.
- The optimal tax equates consumption across agents and it is **progressive**.
- This is the natural benchmark: full information delivers the classic redistributive prescription, with no financial informality.

Benchmark: The Government Observes Income Details

- Suppose (only for this slide) the government observes e directly and taxes $\tau(e)$.
- Since $R > 1$, the couch is strictly dominated: **nobody hides income**.
- The optimal tax equates consumption across agents and it is **progressive**.
- This is the natural benchmark: full information delivers the classic redistributive prescription, with no financial informality.
- **The rest of the talk asks:** what happens once the government loses sight of income, and only sees bank deposits?

Main Model: Taxing Savings, Not Income Assumptions

- Now the government observes *only* bank savings b; tax schedule is $\tau : [0, e_n] \to \mathbb{R}$. Agents who report $b = 0$ (**financially informal**) or split savings across both accounts (**evade**) respond differently to τ.

Main Model: Taxing Savings, Not Income Assumptions

- Now the government observes *only* bank savings b; tax schedule is $\tau : [0, e_n] \to \mathbb{R}$. Agents who report $b = 0$ (**financially informal**) or split savings across both accounts (**evade**) respond differently to τ.

- Agent's problem:

$$\max_{b,\, b^c \geq 0} \quad u\big(e - b - b^c - \tau(b)\big) + \beta\, u\big(Rb + b^c + \psi g\big).$$

- FOC for an interior bank-savings choice:

$$u'\left(c^{young}\right)\left(1 + \tau'(b)\right) = u'\left(c^{old}\right).$$

Main Model: Taxing Savings, Not Income

Assumptions

- Now the government observes *only* bank savings b; tax schedule is $\tau : [0, e_n] \to \mathbb{R}$. Agents who report $b = 0$ (**financially informal**) or split savings across both accounts (**evade**) respond differently to τ.

- Agent's problem:

$$\max_{b,\, b^c \geq 0} \quad u\big(e - b - b^c - \tau(b)\big) + \beta\, u\big(Rb + b^c + \psi g\big).$$

- FOC for an interior bank-savings choice:

$$u'\left(c^{young}\right)\left(1 + \tau'(b)\right) = u'\left(c^{old}\right).$$

- **Key trade-off:** $\tau'(b) \lesseqgtr R - 1$ determines whether the agent prefers the bank or the couch *at the margin*.

$$\text{bank only} \;\Rightarrow\; \tau'(b) < R - 1, \qquad \text{couch only} \;\Rightarrow\; \tau'(0) > R - 1.$$

Two Natural Tax Schedules Both Fail

Decreasing taxes

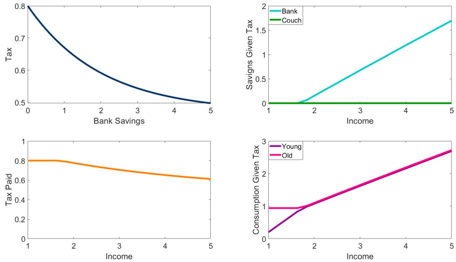

Very regressive: the poor pay the most and the rich save more, worsening inequality.

Progressive taxes

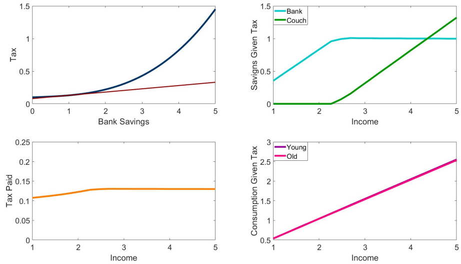

Rich agents hit $\tau'(b) = R - 1$ and **pool**: they save no more in the bank and evade the rest — **evasion at the top**.

Evasion Is Never Optimal

Proposition

Suppose $\tau \in \mathcal{T}$ induces at least one agent e with *both* $b_{\tau,g}(e) > 0$ and $b^c_{\tau,g}(e) > 0$ (an evader) in equilibrium. Then τ cannot be optimal.

- **Why:** an evading agent sits exactly at $\tau'(b) = R - 1$ — the government is taxing away *the entire* marginal benefit of formal saving for that agent.
- Cut the tax rate marginally around that agent's b (so $\hat{\tau}'(b) < R - 1$): she strictly prefers to move couch savings into the bank, raising *her* welfare.
- This local change can always be built to leave every other agent's choice — and public good provision g — unaffected.
- **Punchline:** any tax schedule that induces evasion can always be strictly improved upon. **This is the building block behind "zero evasion" in the main characterization.**

Main Result: There Is An Optimal Concave Tax

Non-Uniqueness

Proposition

Let $\tau^\star$ be an optimal (concave) tax schedule. Then:

1. $\tau^\star$ is non-decreasing.
2. There exists $\tilde{b}$ such that $\tau^\star(b) = \bar{\tau}$ for all $b \geq \tilde{b}$.
3. There is **zero evasion** at $\tau^\star$.
4. Agents with $b \geq \tilde{b}$ consume the same amount young and old (perfect smoothing), but this level differs across them.

What This Looks Like

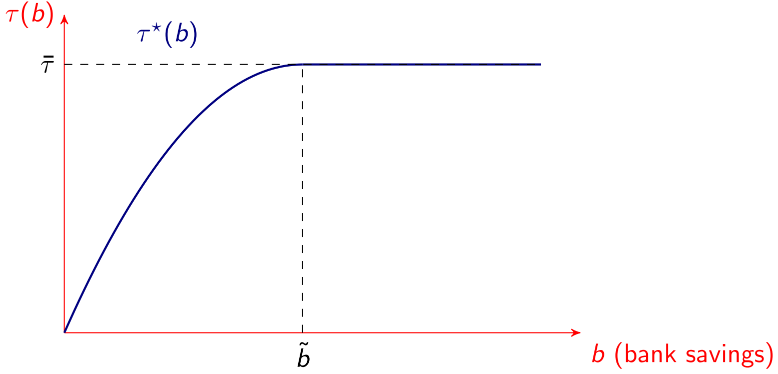

- **Bottom:** taxes rise fast ($\tau'(0) > R - 1$) $\Rightarrow$ the poorest optimally choose the couch (financial informality, but not evasion).
- **Top:** taxes flatten out ($\tau' = 0$) $\Rightarrow$ richer agents keep saving in the bank without losing consumption smoothing, and have no reason to evade.

Why This Shape? Why Not a Step Tax? Why Not Linear?

- **Zero marginal tax rate at the top:** reminiscent of Mirrlees (1971), but for a different reason.
  - Mirrlees: constant taxes avoid discouraging effort at the top.
  - **Here:** the government's ability to tax is bounded above by $R-1$, the return gap between bank and couch. Tax more than that, and the rich walk to the couch.
- By keeping $\tau'(b) < R-1$ in a neighborhood of every equilibrium choice no one wants to split savings between bank and couch.
- Result: **financial informality at the bottom is tolerated**, but **evasion is fully eliminated** everywhere.

Quantitative Illustration: The Income Distribution Matters

Uniform Income

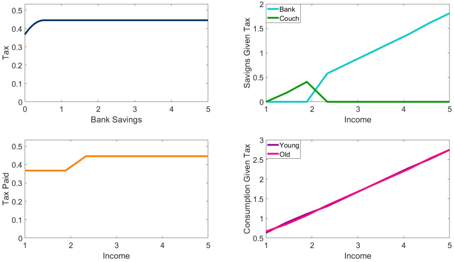

Left-Skewed Income

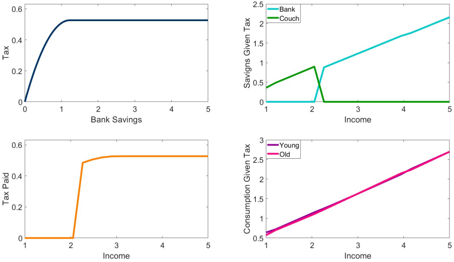

Financial informality rate: **30%** (uniform) vs. **3.6%** (left-skewed, richer economy).

A richer tax base lets the government exempt the poor without sacrificing public good provision.

Public Good Provision and Financial Informality Across Economies

Income Distribution	Interest Rate	Public Good $g^\star$	Fin. Informality Rate
Uniform	$R = 1.15$	0.42	30.0%
Right-Skewed	$R = 1.15$	0.34	66.9%
Left-Skewed	$R = 1.15$	0.49	3.6%
Uniform	$R = 1.05$	0.15	50.0%
Right-Skewed	$R = 1.05$	0.12	84.6%
Left-Skewed	$R = 1.05$	0.14	12.6%

- **Fixing** R**:** a richer income distribution (left-skewed) $\Rightarrow$ more $g^\star$, less financial informality.
- **Fixing the distribution:** a lower R tightens $R - 1$, the government's "tax ceiling" $\Rightarrow$ less revenue, more financial informality.

Policy Takeaways

- **Zero financial informality is not optimal.** Even a benevolent, welfare-maximizing government should tolerate financial informality among the poorest.
- **Financial informality is an endogenous margin, not a policy failure.** It plays the role of a participation constraint that lets the government tax the rich without inducing evasion.
- **Progressive taxes on savings are counterproductive** once agents can escape into financial informality: they push the wealthy out of the formal system.
- **Financial development raises tax capacity.** Since $R - 1$ bounds feasible taxation, policies that raise the return to formal saving expand what the government can collect.

Related Literature Details

- **Optimal taxation under asymmetric information:** Mirrlees (1971), Diamond and Mirrlees (1971).
- **Closest paper, opposite result:** Vairo (2025) also gets asymmetric information + revenue motive, but finds *convex* taxes optimal — because the friction is a robustness concern, not an outside option.
- **Informality:** World Bank (Perry et al. 2007; Loayza and Rigolini 2011), Ulyssea (2020), Bosch and Maloney (2008).
- **Contribution:** a new friction — hidden savings — that overturns the progressivity result once agents can walk away from the tax base.

Conclusion

- When the government can only tax observable savings, taxing savings is **not** the same as taxing income.
- An optimal response is a **concave, eventually flat** tax: it protects redistribution from richer agents while *rationalizing* some financial informality at the bottom.
- This reframes financial informality: not simply a symptom of weak enforcement, but sometimes the **efficient** outcome of an incentive-compatible tax design.
- **Next steps:** general equilibrium interest rates, enforcement/audit technology, political-economy constraints on commitment.

Public Good Provision and Optimal Taxation in a Hidden Savings World

Alberto Ramirez de Aguilar

Universidad Anáhuac

ITAM Alumni Conference, August 2026

Back to Talk

Benchmark Model: Details Back

- Government solves:

$$\max_{\tau:E\to\mathbb{R}} \sum_{k=1}^{n} \Big[u\big(e_k - \tau(e_k) - b(e_k) - b^c(e_k)\big) + \beta u\big(Rb(e_k) + b^c(e_k) + \psi g\big)\Big]\pi_k$$

subject to:

$$\big(b(e_k), b^c(e_k)\big) \in \underset{b,\, b^c \geq 0}{\operatorname{argmax}}\; u\big(e_k - \tau(e_k) - b - b^c\big) + \beta u\big(Rb + b^c + \psi g\big) \;\; \forall\, e_k \in E,$$

$$g = \sum_j \tau(e_j)\pi_j.$$

- Since $R > 1$ strictly dominates the couch, solving each agent's inner problem gives $b^c(e_k) = 0$ and $b(e_k) = \max\left\{\frac{e_k - \tau(e_k) - \psi g}{1+R}, 0\right\}$, which reduces the government's problem to a standard (unconstrained-argmax) optimization over $\{\tau(e_j)\}_j$.
- **Proposition.** The optimal $\tau^\star$: (i) is unique; (ii) equates consumption across agents with positive savings; (iii) is progressive whenever savings are everywhere positive; (iv) implies zero financial informality and zero evasion.

Benchmark: Progressive vs. Non-Progressive Taxes Back

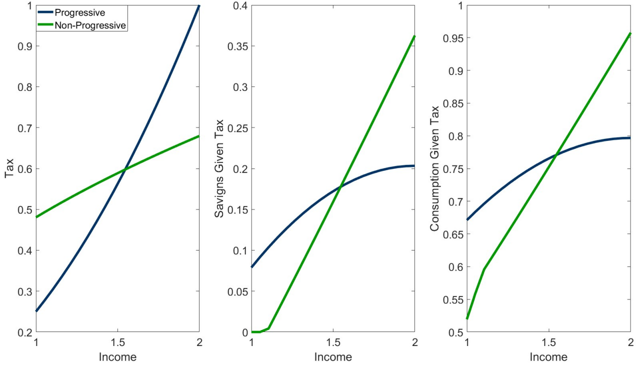

Savings/consumption under a progressive tax $\tau(e) = 0.25e^2$ (blue) vs. a non-progressive tax $\tau(e) = \alpha\sqrt{e}$ (green) delivering the same g. Progressive taxes equalize consumption; non-progressive taxes do not.

Hidden Savings Model: Two Key Assumptions Back

- **Assumption 3 (bounded risk aversion).** With Arrow–Pratt measure $A(c) = -u''(c)/u'(c)$, assume $0 < A(c) < 1$ for all c: risk-averse, but not excessively so.
- **Assumption 4 ($\tau \in \mathcal{T}$).** τ is twice continuously differentiable, bounded, with a uniformly bounded first derivative, and
$$\tau'(b) \geq -1, \qquad \tau''(b) + \left(1 + \tau'(b)\right)^2 \geq 0 \quad \text{for all } b.$$

Hidden Savings Model: Two Key Assumptions Back

- **Assumption 3 (bounded risk aversion).** With Arrow–Pratt measure $A(c) = -u''(c)/u'(c)$, assume $0 < A(c) < 1$ for all c: risk-averse, but not excessively so.
- **Assumption 4 ($\tau \in \mathcal{T}$).** τ is twice continuously differentiable, bounded, with a uniformly bounded first derivative, and

$$\tau'(b) \geq -1, \qquad \tau''(b) + \left(1 + \tau'(b)\right)^2 \geq 0 \quad \text{for all } b.$$

Where does Assumption 4 come from? A Riccati equation.

Let $z(b) := 1 + \tau'(b)$. The condition is exactly the Riccati-type inequality $z'(b) + z(b)^2 \geq 0$. Apply the standard Riccati (logarithmic-derivative) substitution $w(b) := e^{\,b+\tau(b)} > 0$, so $w'(b) = w(b)\, z(b)$. Differentiating once more:

$$w''(b) = w(b)\Big[\underbrace{\tau''(b) + \left(1 + \tau'(b)\right)^2}_{\geq 0 \text{ by Assumption 4}}\Big].$$

Since $w(b) > 0$ always, **Assumption 4** $\iff$ $w(b) = e^{\,b+\tau(b)}$ **is convex** — a mild condition satisfied by lump-sum, linear, flat, progressive, and mildly convex taxes, and exactly what is needed for each agent's problem to be strictly concave.

Hidden Savings: Admissible Tax Schedules $\mathcal{T}$ Back

- $\tau \in \mathcal{T}$: continuously twice differentiable, bounded, uniformly bounded first derivative, and

$$\tau'(b) \geq -1, \qquad \tau''(b) + \big(1 + \tau'(b)\big)^2 \geq 0 \quad \text{for all } b.$$

- These conditions guarantee each agent's problem is strictly concave and differentiable, so the first-order approach is valid (Rogerson, 1985).

- $\mathcal{T}$ still contains lump-sum, linear, flat, progressive, and convex taxes — it only rules out taxes so concave that they break concavity of the agent's problem.

- **Lemma (savings given τ).** If $b_{\tau,g}(e) > 0$, $b^c_{\tau,g}(e) = 0$: $\tau'(b) \leq R - 1$. If $b_{\tau,g}(e) = 0$, $b^c_{\tau,g}(e) > 0$: $\tau'(0) \geq R - 1$. If both positive (evasion): $\tau'(b) = R - 1$ exactly.

Concavification: WLOG Focus on Concave Taxes Back

- **Proposition.** For any optimal $\tilde{\tau}^{\star}$ there exists a concave $\tau^{\star} \in \mathcal{T}$ implementing the same equilibrium outcome.
- Intuition: equilibrium savings choices are *strict global optima*. Non-concave regions of τ away from any agent's chosen b can be replaced by a concave interpolation without changing anyone's incentives.
- This is what lets us restrict attention, without loss, to the class of concave schedules characterized in the main result.

Optimal Tax Is Not Unique Back

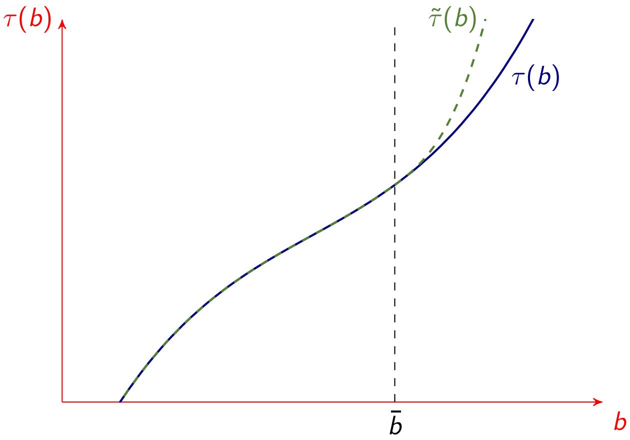

Let $\bar{b}$ be the largest bank savings any agent chooses under τ. Charging *even higher* taxes above $\bar{b}$ changes nobody's behavior (nobody was choosing $b > \bar{b}$ anyway), so $\tilde{\tau}$ is also optimal.

Why Not a Step Tax? Back

A Step Tax Induces Evasion

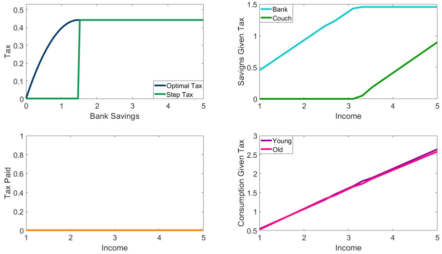

Optimized Steps Converge to $\tau^{\star}$

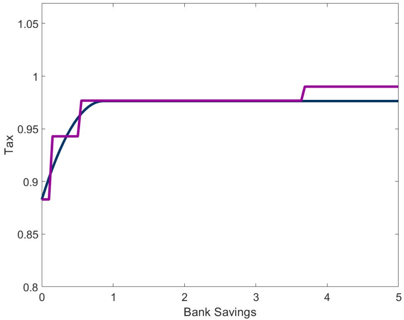

A step tax (zero, then $\tilde{\tau}$) looks simpler, but agents just below $\tilde{b}$ shade their bank savings down to $\hat{b} < \tilde{b}$, pay zero tax, and stash the rest $(b - \hat{b})$ in the couch: **evasion reappears**, and revenue can collapse to zero. But if we instead search for the *best* step tax and let the number of steps grow, the optimized schedule (purple, 5 steps) **converges** to the smooth optimal concave/flat tax $\tau^{\star}$ (blue) — confirming $\mathcal{T}$ is the right class to optimize over.

Why Not a Linear Tax at $\tau' = R - 1$? Back

- **Tempting idea:** set $\tau(b) = (R-1)b$ everywhere. Agents seem indifferent between bank and couch, so "assume" they pick the bank, then collect $(R-1)b$ on everyone.

- **Problem 1 (indeterminacy).** Let $s := Rb + b^c$. Under this τ:

$$c^{young} = e - b - b^c - (R-1)b = e - s, \qquad c^{old} = Rb + b^c + \psi g = s + \psi g.$$

  The split (b, b^c) *never enters the objective* — the agent is indifferent across *every* allocation from all-couch to all-bank, not just the two corners. "Assume they pick the bank" is an unjustified selection; nothing in the model delivers it. (This is exactly why $\tau'(b) = R - 1$ is the evasion boundary in the savings Lemma.)

- **Problem 2 (dominated even if you grant bank selection).** Let b_n be the richest type's chosen bank savings. Flatten τ above b_n: unchanged below, $\tau(b) = \tau(b_n)$ constant above. Revenue is *unchanged* (no one else is affected; type e_n still pays exactly $\tau(b_n)$). But now type e_n faces the undistorted price $1/R$ above b_n and strictly prefers to save more — a **costless Pareto improvement**. So the linear tax cannot be optimal.

Punchline: this is exactly "no distortion at the top." It is why $\tau^{\star\prime} = 0$ must hold for the richest types, not merely $\tau' < R - 1$.

Quantitative Illustration: Right-Skewed Income Back

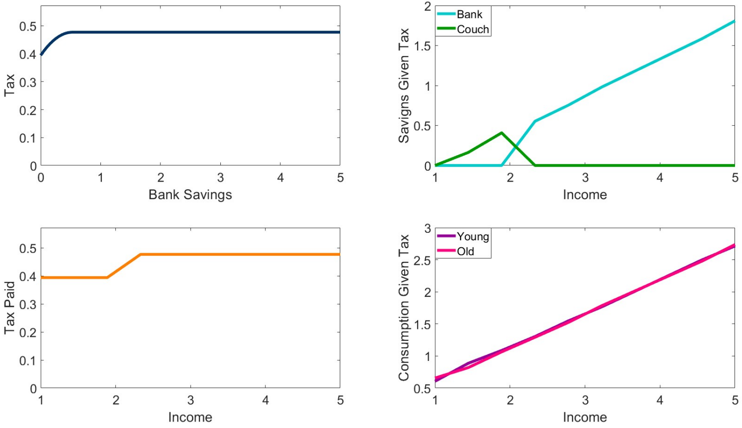

With a large mass of low-income agents, the government cannot afford to exempt them: it must tax broadly and accept a much higher financial informality rate ($\approx 67\%$ at $R = 1.15$) to keep any public good provision.

The Role of the Interest Rate R Back

Low R

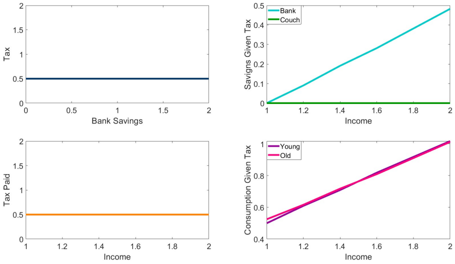

High R

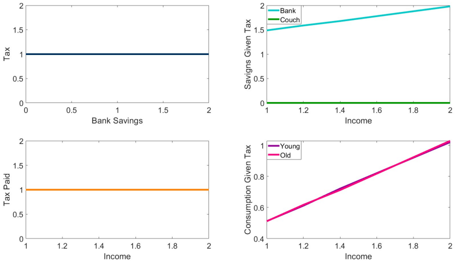

Since $R-1$ bounds the marginal tax rate the government can sustain, a higher bank return mechanically raises the flat tax level $\bar{\tau}$ that agents are willing to pay without defecting to the couch.

Related Literature: Details Back

- **Mirrlees tradition:** Mirrlees (1971), Diamond and Mirrlees (1971) — income taxation under unobserved effort; zero marginal tax at the top.

- **Dynamic private information:** Albanesi and Sleet (2006), Golosov and Tsyvinski (2006), Kocherlakota (2010), Farhi and Gabaix (2020).

- **Vairo (2025):** government maximizes welfare + revenue under *uncertainty* about the income process $\Rightarrow$ optimal tax is **convex** (aligns risk-averse government with risk-loving consumers). My model instead has an outside option that *caps* feasible taxation, delivering the opposite (concave) prescription.

- **Informality:** Perry et al. (2007), Oviedo and Mark (2009), Bosch and Maloney (2008), Ulyssea (2020), Meghir et al. (2015), Olken and Singhal (2011).

Timing and Equilibrium Definition Back

- **Timing:** (1) government commits to τ; (2) nature assigns income types $e \sim F$; (3) agents choose (b, b^c) given τ, g.

- This commitment timing makes the problem a **screening** problem (as in mechanism design), not a signaling one.

- **Equilibrium** $\left(g, \{(b_{\tau,g}(e), b^c_{\tau,g}(e))\}_{e \in E}\right)$: each agent optimizes given (τ, g), and the budget balances, $g = \sum_i \tau(b_{\tau,g}(e_i))\pi_i$.

- For $\tau \in \mathcal{T}$, an equilibrium g exists and the correspondence $\mathcal{G}(\tau)$ is in fact a continuous function.