Public Good Provision and Optimal Taxation in a Hidden Savings World*

Alberto Ramirez de Aguilar†

August 2026

Abstract

This paper studies the optimal design of tax policy and public good provision in an economy where individual income is unobservable, but formal financial transactions are observable. In this environment, agents choose how much of their income to save either in a taxed, observable bank account or in an untaxed, hidden account (their "couch"). I show that financial informality and tax evasion arise endogenously as responses to the tax schedule, and I characterize the welfare-maximizing tax structure under these considerations. I characterize the optimal tax schedule within the class of concave tax schedules, showing that it is flat beyond a threshold: it deters evasion at the top, and induces agents to save in the formal sector without taxing away the benefits of doing so. I also show that progressive tax schedules become counterproductive once they redistribute enough to tax away the full return on formal saving, since this induces costly evasion. The model suggests that the optimal policy is to rationalize and tolerate some level of financial informality among the lower-income households to sustain incentive compatibility for richer agents and prevent evasion. The results highlight a trade-off between redistribution and public goods provision in the presence of hidden savings and offer novel insights into tax design for economies where financial informality is widespread.

*I am grateful to my committee, Harold L. Cole, George J. Mailath, and Alessandro Dovis, for all the feedback and support I have received for this and other projects. I would also like to thank Dirk Krueger, Joachim Hubmer, Lukas Nord, Eddie Dekel, and all the participants of the Penn Macro Group and the Penn Theory Group for their helpful comments.

†Universidad Anáhuac México. E-mail: alberto.ramirezag@anahuac.mx.

1 Introduction

How should a government design tax policy and provide public goods when it cannot observe individual incomes but can monitor financial activity? This question is particularly pressing in economies where a large share of household savings and financial activity takes place outside the formal banking system, so that income-based taxation becomes infeasible and governments must instead rely on observable financial proxies, such as bank deposits, to collect revenue. This creates an incentive for agents to conceal income through informal, off-the-books savings, undermining both redistribution and revenue collection.

This paper develops a framework in which agents can choose to save either in a taxed, observable bank account or in an untaxed hidden account. I study the optimal tax schedule when the government cannot observe income but can condition taxes on bank savings. In this environment, financial informality and evasion arise endogenously: agents may avoid the financial system altogether or use it strategically to reduce their tax burden. Throughout the paper, I use the term *financial informality* to refer specifically to this savings-concealment margin: an agent who forgoes the (taxed) banking system entirely in favor of an untaxed store of value. This is a narrower notion than the informality studied in the labor- and firm-informality literatures, which concern unregistered employment and unregistered businesses; here, agents' income and labor-market status are exogenous and unaffected by the tax schedule, and the only margin of adjustment is where, not whether, to save. I characterize the welfare-maximizing tax schedule, taking into account the effects of taxation on savings behavior and the government's capacity to fund public goods.

The framework is a simple two-period version of the classic public good contribution game. Each agent in the economy is endowed with an income in the first period (when young) and must decide how much to save for the second period (when old). Agents can choose to save in either a bank or their "couch", the bank paying a larger gross return than the couch. Income when old consists of savings when young plus the provision of a public good (interpreted as healthcare system). To provide the public good, the government must collect taxes from agents when young.

To isolate the role of hidden savings and asymmetric information, I first analyze a benchmark model in which the government observes agents' incomes and can condition taxes directly on them, with the objective of maximizing utilitarian welfare. The optimal tax is progressive, allowing the government to redistribute across income types without inducing agents to hide savings. Since the taxed bank account earns interest while the untaxed account does not, agents always prefer to save formally, and financial informality does not arise. In this case, the government achieves redistribution, efficient savings behavior, and a large public good provision. Importantly, in this frictionless setting, agents have no incentive to hide income, and hence have no incentive to use the hidden account (since it pays a lower interest on savings).

The main contribution of the paper is to show how this result changes when the government cannot observe income and must rely on taxed financial activity as the basis for taxation. In this setting, agents may respond to taxation by hiding part or all of their income in informal savings. The presence of an outside option imposes a new incentive constraint: the tax schedule must discourage financial informality and evasion, which limits the government's ability to redistribute. I characterize the optimal tax schedule within the class of concave tax schedules, showing that it is flat beyond a threshold.[1] This structure keeps richer agents in the formal financial system, avoids evasion at the top, and accommodates financial informality among some lower-income households. Even though the government values both redistribution and revenue (through public good provision), a tax schedule that is progressive enough to tax away the full return on formal saving cannot be optimal, since doing so induces costly evasion. The outside option thus shapes the optimal policy and limits the feasibility of steeply progressive taxation, within the class of concave schedules for which I am able to provide a full characterization.

2 Related Literature

This paper contributes to three strands of the literature. First, it builds on the Mirrlees tradition of optimal taxation under asymmetric information, extending it to account for financial informality and endogenous evasion. Second, it offers a theoretical complement to empirical work on informality, highlighting how tax policy can shape financial-sector participation. Third, it provides normative guidance for economies where widespread financial informality and limited enforcement capacity make the trade-offs between redistribution, tax collection, and public goods—in the classic Samuelsonian sense (Samuelson, 1954)—particularly acute.

This paper fits within the extensive literature on optimal taxation and asymmetric information. Classical taxation models such as Mirrlees (1971) analyze optimal income taxation, where individual productivity is private information and the government conditions taxation on observable income. In this class of settings, the government seeks to collect taxes because it needs to finance an exogenously given project and has the utilitarian welfare utility function as an objective. This model can be embedded within the classic principal-agency models in which the government (principal) seeks a rule (taxes) that maximizes an objective function (utilitarian welfare) subject to incentive-compatibility for the agent's choices (individuals). A classic result under suitable conditions is a zero marginal tax rate for the highest type. In classic Mirrleesian-type economies, providing general statements about the optimal tax schedule's overall shape is difficult since marginal tax rates are usually U-shaped: numerically optimal schedules tend to feature high marginal rates at the bottom of the income distribution, where screening low types from mimicking is most costly, declining rates over the middle range, and rates that can rise again before typically falling to zero at the very top of a bounded type space. Because this shape is a numerical regularity rather than a result that follows transparently from primitives, it does not translate into a simple, portable qualitative characterization (e.g., globally progressive or regressive) of the kind I am able to obtain in this paper's hidden-savings environment.

My paper can be thought of as a Mirrleesian economy, but in which households have an outside option, call it house production. Then, the income that the government observes, which is given by the income generated by non-household production, may be underestimating the household's entire income, which I interpret as evasion. As I highlight in my main results, this outside option and the possibility to evade taxation restrict the capacity of the government to tax and have sharp implications on the shape of optimal taxes: within the class of concave schedules, for which I provide a full characterization, the optimum is flat beyond a threshold

1 I do not prove that this restriction is without loss of generality relative to the unconstrained problem. As a numerical check, I solved the government's problem without imposing concavity and found that the resulting optimal schedule was concave in every parameterization considered; see Section 5 for details.

rather than progressive.

A recent paper by Vairo (2025) also studies optimal income taxation under asymmetric information but reaches a contrasting result: the optimal tax schedule is convex (i.e., progressive). The key distinction lies in the government's objective. In Vairo (2025)'s framework, the government seeks to maximize a combination of utilitarian welfare and tax revenue and is ambiguity-averse (in the max-min sense) about the income-generating process of agents. In this setting, a concave tax schedule creates a misalignment between the government, which is risk-averse due to the concavity of revenue, and consumers, who become effectively risk-loving as their consumption becomes convex. As a result, convex taxes are optimal since they induce concave consumption and align incentives. My paper, while featuring asymmetric information and a government that cares about welfare and revenue (through public good provision), reaches the opposite conclusion: the optimal tax schedule I characterize, within the class of concave schedules, is concave rather than convex. This is because agents in my model have an outside option (financially informal savings) that they can turn to if taxes become too progressive. This outside option constrains the tax authority: a tax schedule that is progressive enough to tax away the full return on formal saving induces costly evasion and reduces public good provision, which is why the concave characterization I obtain is not itself progressive. In this sense, my results align more with the Mirrlees tradition, emphasizing how hidden income and evasion concerns fundamentally reshape the set of implementable and welfare-improving tax schedules.

My paper is also closely related to the literature on efficient allocations and optimal taxation when agents can privately store resources. Cole and Kocherlakota (2001) study environments in which agents have private income and access to a hidden storage technology, and show that the mere availability of hidden storage caps the insurance a planner can provide, since agents can always fall back on self-insurance outside the mechanism. My "couch" technology plays an analogous role: it is a dominated but hidden store of value that agents turn to whenever bank savings are taxed too aggressively, which is precisely why the optimal tax schedule must flatten out rather than keep rising. This also connects to the classic finding, formalized by Golosov et al. (2003) and Kocherlakota (2005) and surveyed in Kocherlakota (2010), that savings should generally be distorted away from their efficient level under private information (the inverse Euler equation). I reach a related conclusion through a different instrument: rather than an implicit wedge on an unobserved consumption-savings margin, the distortion here is an explicit, observable tax schedule on bank savings, whose steepness is bounded by the return on the competing hidden technology. Albanesi and Sleet (2006) and Farhi and Gabaix (2020) study other departures from the classic Mirrleesian benchmark, through intertemporal distortions and behavioral frictions respectively, further illustrating how sensitive the shape of optimal taxation is to the frictions assumed. My paper contributes to this literature by introducing a new friction, hidden savings, that interacts with the government's redistributive and revenue goals in a novel way, leading to new insights about the feasibility and desirability of progressive taxation.

The mechanisms I study connect to a body of empirical evidence on the scale and consequences of financial informality and tax-driven evasion. Gordon and Li (2009) argue that a defining feature of tax systems in developing countries is that agents can largely escape taxation simply by remaining outside the banked financial sector, which governments rely on as their primary source of information on economic activity. Chodorow-Reich et al. (2020) provide direct evidence of this margin at work, showing that a substantial share of economic activity in India was conducted in cash outside the banking system prior to the country's 2016 demonetization, and that districts more exposed to the reform experienced sharp, if temporary, contractions in output and bank credit as this activity was pushed through the formal financial system. Alstadsæter et al. (2019) use leaked data from offshore financial institutions to show that tax evasion through undisclosed accounts is heavily concentrated among the wealthiest households, who evade a far larger share of their tax liability than random audits alone would suggest. Taken together, this evidence points to a hidden, untaxed store of value—whether cash, informal channels, or offshore accounts—as a first-order determinant of both tax evasion and the shape of feasible tax policy. My framework offers one possible rationalization of these patterns: financial informality arises endogenously as agents' equilibrium response to a tax schedule that would otherwise tax away too much of the return to formal savings, and the optimal tax schedule I characterize is precisely one that accommodates this margin rather than assume it away.

The remainder of the paper is structured as follows. Section 3 presents the model environment. Section 4 analyzes the benchmark case with observable income. Section 5 introduces the hidden savings setting and characterizes optimal behavior and taxation under these constraints. Section 6 discusses the quantitative implications of the model, particularly about how the primitives of the model shape the optimal tax schedule, and provides a cross-country illustration of the model's central mechanism. Section 7 discusses how the paper's main results are affected by several of the model's simplifying assumptions.

3 Model Setup

I consider an economy with a continuum of agents, represented by the $[0, 1]$ interval. Each agent is characterized by an exogenously given type e, representing the agent's income. Income has a finite support $E = \{e_1, ..., e_n\}$ with $e_1 < e_2 < ... < e_n$. The probability that agent $i \in [0, 1]$ has income e_j is given by $0 < \pi_j < 1$. I denote F to be the distribution associated with the type assignment. Upon birth, each agent receives an income drawn independently and identically from F.[2]

In this economy, agents live for two periods: in their first period, which I refer to as the "young" stage of the agent's life, agents receive an exogenously determined income, must pay taxes, and decide how much to save for the future. In the second period, which I refer to as the "old" stage, agents consume using their savings and the provision of a public good. As I describe below, the provision of the public good will be a function of the aggregate tax collection done to young agents.

When agents are young, they have two ways of saving: they can deposit in a bank, which pays a gross interest of $R > 1$; or they can save in their "couch," which pays a gross interest of one.[3]

2 As is standard in continuum-of-agents models, I invoke the exact law of large numbers to treat the realized cross-sectional share of each type e_j in the population as the deterministic constant π_j, rather than as a random variable subject to sampling variation; see, e.g., Judd (1985) and Uhlig (1996) for a discussion of this convention and the measure-theoretic subtleties involved in applying the law of large numbers to a continuum of independent draws. This is what justifies treating aggregate tax revenue and public good provision, both of which are population-weighted sums over $\{\pi_j\}$, as non-random throughout the paper.

3 The value one is a normalization, not a substantive restriction. The couch is the untaxed outside option against which formal savings are measured, and what the model requires throughout is only that the bank offer a strictly higher return than this informal alternative, i.e., a positive spread between the two returns. Setting the couch's return to one is simply a choice of units, so R should be read as the bank's return *relative to* the couch, rather than as an interest rate in some absolute or economy-wide sense.

Both types of savings are paid together with their corresponding interest when agents become old.

I am interested in characterizing the optimal tax schedule in this environment and the implied individual consumption and savings behavior. I study this problem considering two scenarios:

1. Benchmark: the government can perfectly observe the income of every agent in the economy.
2. The government does not know the income of any agent. However, it can observe how much the agent saves in the bank.

In both cases, I assume a benevolent government aiming to maximize aggregate welfare. Importantly, the following sections perform a "partial equilibrium" analysis since I characterize the economy's optimal taxes, consumption, and savings behavior, taking the interest rate R as a parameter and not as an endogenously determined equilibrium object. This simplifies the analysis substantially and also allows me to perform comparative statics on how optimal taxes change as a function of R.[4]

One particular type of tax that may sound appealing from a redistributive perspective is a progressive tax. This type of tax implies that richer agents end up paying more. One of the objectives of this paper is to understand how hidden information may affect the plausibility of progressive taxation. I now define what I mean by progressive taxes throughout this paper.

Definition 1. Given a tax schedule τ, the marginal tax rate for $e \in \{e_2, ...e_n\}$ is:

$$MTR(e_i) = \frac{\tau(e_i) - \tau(e_{i-1})}{e_i - e_{i-1}}.$$

A tax schedule is **progressive** if the marginal tax rate is non-decreasing in e_i.

The main focus of this paper is to examine how the tax schedule may be pushing agents in the economy towards financial informality/evasion and to see if the welfare-maximizing tax implies no financial informality (and hence from a normative point of view, financial informality could be thought of as "bad" for the economy's welfare). Throughout, *financial informality* refers narrowly to an agent's choice to conceal savings from the formal banking system; it is not a statement about the agent's labor-market status or the legal status of her employer, both of which are exogenous in this model. In the context of my model, I define an agent to be **financially informal** if she does not use the financial system of the economy at all, while an agent is an **evader** if she has both positive bank savings and positive couch savings. In my main model, the government uses bank savings to determine tax payments. Hence, if an agent is not using the bank, she is completely concealing her income from the government (which is why I call this agent financially informal). In contrast, if the agent does save but does not completely "report" her entire savings, the agent is an evader.

4 Section 7 discusses how the paper's results are affected by relaxing the assumption that R is exogenously fixed rather than determined in general equilibrium.

4 Benchmark Model

In this framework, the government can perfectly observe the income of each agent. Hence, the government can condition the tax collection of agents on their type. In this sense, the government chooses a tax schedule $\tau : E \to \mathbb{R}$ to maximize welfare in the economy. Notice that the range of this tax schedule allows for subsidizing some of the agents.

In this world, the utility of an agent of income e_k is:

$$u\left(e_k - \tau(e_k) - b - b^c\right) + \beta u\left(Rb + b^c + \psi g\right),$$

where $b, b^c \geq 0$ represent the amount of bank/couch savings; $\tau(e_k)$ are the taxes that an individual of type e_k must pay; g is the amount of public good provided to all agents; $\beta > 0$ represents the agent's patience; and $\psi > 1$ is a parameter that captures the positive externality of public good provision. Importantly, when each agent decides how much to consume (and therefore save) in each period, they take the tax schedule and g as given. Throughout the paper, I focus on the case $R < \psi$, which guarantees that a positive amount of public good is provided in equilibrium, and therefore, collecting taxes is a non-trivial problem.[5]

The provision of the public good is determined by the tax collection (net of any potential subsidies):

$$g = \sum_{j=1}^{n} \tau(e_j)\pi_j.$$

Timing in this framework is as follows: first, the government commits to a tax schedule $\tau : E \to \mathbb{R}$ seeking to maximize welfare; then nature assigns each agent an income type according to F; lastly, upon observing τ and g, each agent chooses how much to consume and save to maximize their utility. Importantly, the government cannot modify ex-post its tax rule, which is a reasonable assumption since usually a country has to go through lengthy reforms to modify taxes, and agents are aware of these modifications before they are enforced. I maintain this commitment assumption throughout the paper.

In this sense, the government solves the following problem:

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

$$(b(e_k), b^c(e_k)) \in \text{argmax}_{b, b^c \geq 0}\ u\left(e_k - \tau(e_k) - b - b^c\right) + \beta u\left(Rb + b^c + \psi g\right) \ \ \forall\ e_k \in E,$$

$$g = \sum_{j=1}^{n} \tau(e_j)\pi_j.$$

I now move on to analyze each agent's problem, taking as given both the tax schedule and the public good provision.

5 The strict inequality also rules out a degenerate case at $\psi = R$: if $\psi = R$, a common shift $\tau(e_k) \mapsto \tau(e_k) + s$ applied to every type raises g by s but leaves every positive-saver's consumption when young and old exactly unchanged (see the proof of Proposition 1, part 1, in Appendix A), so the welfare-maximizing tax schedule would fail to be unique whenever all agents have positive savings. Requiring $\psi > R$ avoids this.

4.1 Savings and Consumption Decisions Given Taxes

Throughout this subsection, let τ and g be fixed. I make the following two assumptions throughout the paper.

Assumption 1. $\beta R = 1$.

Assumption 2. $u(\cdot)$ is a strictly increasing, twice continuously differentiable, and strictly concave function.

Both are fairly standard assumptions made in many macroeconomic models. The first one is assuming that the market interest rate is equal to the agent's discount rate, which induces perfect consumption smoothing over time (under non-distortionary taxes); while the second assumption implies a decreasing marginal benefit of consumption, which translates into demands having the standard property of being decreasing with respect to prices (in this case, the interest rate).

Given these assumptions, the problem that a type e_k agent solves,

$$\max_{b,b^c \geq 0} \; u\left(e_k - \tau(e_k) - b - b^c\right) + \beta u\left(Rb + b^c + \psi g\right),$$

has a unique solution, and the first-order conditions are necessary and sufficient to characterize the optimal solution (since this is a strictly concave problem). Moreover, since $R > 1$, optimal couch savings are zero, since any dollar that the agent could save in the couch is dominated by saving it in the bank and receiving a return higher than one. Using the FOCs to characterize the optimal solution, we obtain:

$$b(e_k) = \max\left\{\frac{e_k - \tau(e_k) - \psi g}{1+R}, 0\right\}.$$

Here, we see that the value of taxes and public good provision is important for determining the amount of savings. Another important observation, which becomes relevant in subsequent sections, is that, in this full-information model, the curvature of the tax schedule is irrelevant to determining individual saving/consumption choices. The only relevant value for agent e_k is $\tau(e_k)$, but how the tax changes across different agents is not important to determine individual decisions. Of course, the curvature and level of the tax schedule will be important to pin down the aggregate public good provision, but since agent e_k takes g as given, this does not affect its individual choices.

Figure I displays the implied savings and consumption decisions for each agent given two different tax schedules: one that is progressive (shown in blue) and one that is not (a concave tax structure, shown in green). Since $\beta R = 1$, agents with positive savings consume the same amount when young and old; however, this amount may vary across agents with different incomes. This figure illustrates why progressive taxes may be more "ideal" from an egalitarian perspective, since they imply more equal consumption across agents, whereas consumption variance under non-progressive taxes is higher.

The following lemma summarizes agents' behavior given the tax schedule and public good provision.

Figure I: SAVINGS AND CONSUMPTION DECISIONS GIVEN TAXES.

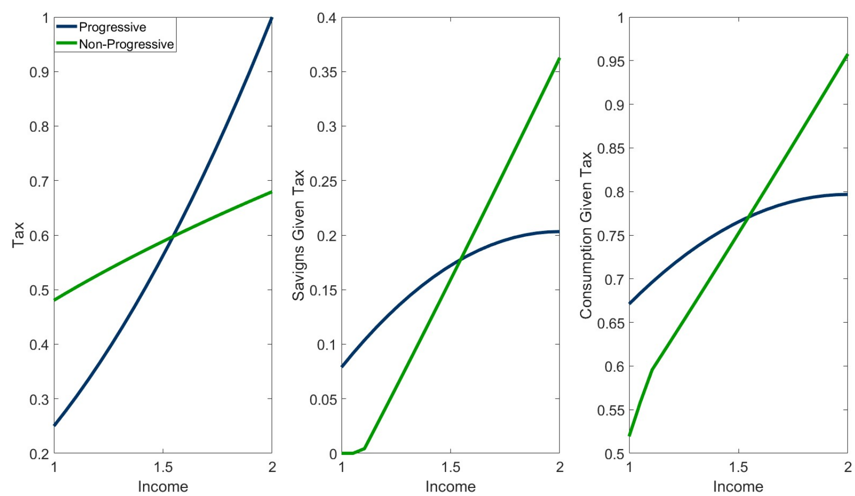

NOTES: This graph plots the optimal savings and consumption (when young and old, which are equal) decisions for two different tax schedules: one progressive ($\tau(e) = 0.25e^2$), displayed in blue, and one that is not progressive ($\tau(e) = \alpha\sqrt{(e)}$), displayed in green. The value of α is such that both tax schedules imply the same public good provision. For this figure, I consider $R = 1.05$, $\psi = 1.2$, and $E \sim U([1, 2])$.

Lemma 1. *Let $\tau : E \to \mathbb{R}$ be fixed as well as g. Then:*

1. *There is a unique $(b(e_k), 0)$ that solves agent e_k's problem.*
2. *If agent e_k has positive savings, then they are decreasing in $\tau(e_k)$ as well as in g.*
3. *Every agent with positive savings consumes the same amount when young and old (although this may vary across agents).*

The proof of Lemma 1 can be found in Appendix A. Intuitively, the first part of this lemma states that every individual saves zero on the couch, no matter the tax schedule they are facing. This happens since the bank has a higher return than the couch ($R > 1$). On the other hand, as taxes are higher, agents have less income when young, and therefore can save less for the future, which is why savings are a decreasing function of taxes. On the other hand, as agents face a larger public good provision g when old, they are less in need to save for the future, and hence their savings also decrease with g. Finally, since $\beta R = 1$ then each agent's discount rate is equal to the interest rate, which implies as in any standard intertemporal model that consumption is perfectly smoothed throughout periods.

4.2 Welfare-Maximizing Tax Schedule

Considering the insights of Lemma 1, I can re-write the government's problem that characterizes the optimal tax schedule as follows:

$$\max_{\{\tau(e_j)\}_j} \sum_{k=1}^{n} \left[u\left(e_k - \tau(e_k) - b(e_k)\right) + \beta u\left(Rb(e_k) + \psi g\right) \right] \pi_k, \quad \text{subject to:}$$

$$b(e_k) = \max\left\{\frac{e_k - \tau(e_k) - \psi g}{1+R}, 0\right\}.$$

$$g = \sum_{j=1}^{n} \tau(e_j)\pi_j.$$

This reformulation is considering the optimal $b(e_k)$ (which can be obtained using the first-order conditions of each agent's problem). In addition, since E is a finite set, and the tax schedule in this problem has a domain equal to E, then finding a tax function to maximize welfare is the same as finding $\{\tau(e_j)\}_j$, with the advantage that formulating the problem this way, turns the government's problem into a standard optimization problem.

The following proposition characterizes the optimal tax schedule in the benchmark model, and its proof can be found in Appendix A.

Proposition 1. *The optimal tax schedule $\tau^\star$ has the following properties.*

1. *It is unique.*
2. *It equates consumption across all agents with positive savings, in the sense that agents with positive savings consume the same when young and old, and this consumption level is the same across all agents with positive savings.*
3. *If all agents have positive savings given $\tau^\star$, then the optimal tax schedule $\tau^\star$ is progressive.*
4. *It implies zero financial informality and zero evasion.*

This proposition states that the optimal tax schedule $\tau^\star$ implies the same consumption across different agents in the economy, which means that the optimal taxes should promote perfect equality in the economy. The third part of the proposition states that the way to achieve this is through progressive taxation. Something to note is that equal consumption is optimal regardless of the particular details on the income distribution F.[6]

The last statement of this proposition says that the optimal tax schedule induces agents not to be financially informal (i.e., have zero bank savings and positive couch savings) nor to evade (have both positive bank and couch savings). This is because bank savings are more attractive for consumers regardless of the shape of the tax schedule.

5 Hidden Savings Model

Now, I turn to analyze the main model of this paper, in which I assume that the government cannot directly observe the income of any agent. The government is only able to observe the amount of savings each agent deposits in the bank (but couch savings are non-observable as well). Therefore, the tax schedule will be a function $\tau : [0, e_n] \to \mathbb{R}$, where the domain represents the bank savings of each agent. Agents cannot borrow and the maximum amount an

6 However, the optimal value of consumption across agents does depend on $\{\pi_j\}_j$.

agent could (potentially) save is e_n.[7]

The motivation behind the assumption that the government cannot observe either the couch savings or the agent's income comes from the fact that in economies with widespread financial informality, many transactions are done using cash. There, the tax authority cannot have precise control of these transactions. The government of these economies can only condition taxes on observables, such as savings in the financial system.

Timing in this framework is as follows: first, the government commits to a tax schedule τ, seeking to maximize the economy's utilitarian welfare;[8] second, nature assigns each agent an income type according to the distribution F; and finally, each agent decides how much to save (in both the bank and couch) to maximize their utility taking as given both the tax schedule τ as well as the provision of the public good g.

Hence, the utility of an agent of type $e \in E$, who takes as given τ, g, is then:

$$u\left(e - b - b^c - \tau(b)\right) + \beta u\left(Rb + b^c + \psi g\right),$$

where b, b^c are restricted to be non-negative.[9] Then, different levels of bank savings result in (potentially) different taxes that the agent must pay, which provides an incentive towards not saving in the bank too much whenever taxes are too high. In this world, the government will have to develop a tax schedule that encourages agents to save in the bank, since otherwise tax collection may be too low (or even zero), and public good provision and redistribution may suffer.

In this framework, I consider the following definition of financial informality and evasion.

Definition 2. An agent $e \in E$ is **financially informal** if, given taxes and public good provision, $b_{\tau,g}(e) = 0$ and $b^c_{\tau,g}(e) > 0$; while an agent is an **evader** whenever $b_{\tau,g}(e) > 0$ and $b^c_{\tau,g}(e) > 0$.

The idea behind these definitions is that agents who do not use the financial system at all and rely on the couch as a mean of savings are those who seek to completely conceal their income from the government, thereby making them financially informal. On the other hand, if an agent uses the bank and the couch in a way that partially conceals its income, it is, in effect, an evader.

In this framework, the curvature of τ is relevant to determine the optimal amount of savings since the agents will internalize how marginal changes in b affect the taxes the agent will have to pay. In particular, the first-order condition of each agent's problem (assuming that taxes are differentiable) with respect to bank savings (assuming a positive optimal solution) is:

$$u'\left(e - b - b^c - \tau(b)\right)\left(1 + \tau'(b)\right) = u'\left(Rb + b^c + \psi g\right).$$

7 Allowing agents to borrow against future income would require specifying a lending counterparty, which the partial-equilibrium approach adopted here (Section 3) does not model: since R is taken as an exogenous spread rather than an equilibrium object, there is no residual claimant on the other side of the loan within the model. A natural extension would let agents borrow at a rate no lower than R, ruling out a pure arbitrage against bank savings, and would relax the constraint $b \geq 0$ without altering the paper's central mechanism, the tradeoff between taxed bank savings and the untaxed couch. I abstract from this margin and leave a full treatment of borrowing to future work.

8 The timing assumed in this model, makes this a "screening problem" which is more aligned with the classic mechanism design literature, like Myerson (1981). Removing the commitment assumption would make this a "signaling problem, " which is interesting but has many other features that are beyond the scope of the paper. A paper that illustrates the distinction between these type of problems is Stiglitz (1987).

9 Unlike the bank savings, $b^c \geq 0$ is a natural restriction since couches do not borrow.

This first-order condition suggests that taxing savings is distortionary and that, in general, taxing savings when income is not observable is not equivalent to taxing income. Moreover, to fully characterize behavior using this equation (i.e., the first-order conditions being necessary and sufficient), the shape of taxes becomes relevant: for this problem to be strictly concave, I not only need to impose some structure in the utility function, but also on how taxes behave with respect to savings. Then, in order to be able to use the classic approach of characterizing the optimal behavior of agents using first-order conditions, I need to impose certain restrictions on taxes. First, I need to assume that τ is differentiable. Moreover, for the agent's problem to be strictly concave (which guarantees that the first-order conditions are necessary and sufficient), more structure must be imposed on τ. I make the following assumptions that are considered throughout this section.

The first assumption concerns the concavity of the utility function and, in particular, the agent's attitude toward risk. The Arrow-Pratt measure of absolute risk aversion for a utility function u at consumption level c is defined as:

$$A(c) = -\frac{u''(c)}{u'(c)}.$$

If $A(c)$ is increasing, the curvature of u becomes more pronounced at higher levels of consumption, indicating that the agent is becoming increasingly risk-averse. The following assumption imposes an upper and lower bound on the values that $A(c)$ can take.

Assumption 3. The utility function $u : \mathcal{D} \to \mathbb{R}$ satisfies $0 < A(c) < 1$ for all $c \in \mathcal{D}$, where $\mathcal{D}$ is the domain set for consumption.

This condition restricts the behavior of the first and second derivatives of u, ensuring that the agent is risk-averse but not excessively so. While it rules out some strictly concave utility functions, many specifications commonly used in the literature comply with this restriction. For instance, the CRRA utility function $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ with $\sigma > 0$ and $\mathcal{D} = [c_0, e_n]$ satisfies the condition as long as $c_0 > \sigma$, which, economically, requires the risk aversion coefficient of the utility function to be small.[10] The same holds for logarithmic utility, $u(c) = \log(c)$, the limiting case of CRRA utility as $\sigma \to 1$: the condition again requires $c_0 > 1$.[11]

Assumption 4. Fix constants $M, M', M'' < \infty$, common to every tax schedule considered in the paper. I consider functions $\tau : [0, e_n] \to \mathbb{R}$ that are continuously twice differentiable and satisfy:

$$|\tau(b)| \leq M \quad \text{for all } b \in [0, e_n],$$
$$-1 \leq \tau'(b) \leq M' \quad \text{for all } b \in [0, e_n],$$
$$\tau''(b) + \big(1 + \tau'(b)\big)^2 \geq 0 \quad \text{for all } b \in [0, e_n],$$
$$|\tau''(b)| \leq M'' \quad \text{for all } b \in [0, e_n].$$

10 This follows from the fact that for CRRA utility, $A(c) = \sigma/c$.

11 Since $A(c) = 1/c$ for log utility, $0 < A(c) < 1$ requires $c > 1$, consistent with the CRRA formula $A(c) = \sigma/c$ evaluated at $\sigma = 1$.

I denote $\mathcal{T}$ as the set of functions satisfying these conditions for the fixed triple (M, M', M''). This is a subset of $\mathcal{C}_2([0, e_n])$, the set of continuous and twice continuously differentiable functions with domain $[0, e_n]$. This, together with the $\mathcal{C}_2$-norm, forms a metric space.[12] I endow $\mathcal{T}$ with the $\mathcal{C}_2$-norm to measure the distance between functions on this set.

The conditions restricting the tax functions τ to belong to $\mathcal{T}$ are sufficient to guarantee that each agent's problem has a unique solution characterized by its first-order conditions.[13] Appendix B provides details on this. Restricting attention to $\tau \in \mathcal{T}$ is a technical assumption allowing me to characterize the optimal solution using first-order conditions and the first-order approach. And, although $\mathcal{T}$ does exclude some twice-continuously differentiable functions, it contains most of the tax schedules that have been of interest to the literature. The following "classic" tax schedules belong to this set: lump-sum taxes, linear taxes, flat taxes (affine taxes), progressive taxes, and convex taxes, provided M, M', and M'' are chosen generously enough to accommodate their levels and slopes on $[0, e_n]$ (e.g. $M' \geq 1$ already covers any marginal rate of realistic interest, and M need only exceed the largest tax level considered). The main restriction that $\mathcal{T}$ imposes is to discard functions whose concavity (and hence $\tau''(b) < 0$) is severe enough to overturn strict concavity of the objective along the bank-savings direction alone. The additional bound $|\tau''(b)| \leq M''$ is not restrictive in any economically relevant sense: every schedule just listed has constant (indeed, zero) curvature, so the bound is satisfied for any $M'' \geq 0$; it is included, together with the now-uniform bounds M and M', solely because a curvature bound and first-derivative bound that hold uniformly across all of $\mathcal{T}$ (not just for each τ individually) are what allow me to establish, in Appendix B, that an optimal tax schedule exists in the first place.

To show that taxing savings when income is unobservable poses a challenge for the government, I now argue that the optimal tax schedule I described in the benchmark model is not implementable in this world. More specifically, there is no tax schedule that induces the same level of savings for each agent and provision of the public good as in the benchmark framework. For simplicity, let's assume taxes are continuous.[14] Also, let $\hat{g}$ be the amount of public good provided under τ. If this same tax schedule were to be implemented under hidden information, then $\hat{g} = \tau(\hat{b})$, since all agents have the same savings level. For $\hat{b}$ to be an interior optimum for an agent of income e, her first-order condition must hold:

$$u'\left(e - \hat{b} - \tau(\hat{b})\right)\left(1 + \tau'(\hat{b})\right) = u'\left(R\hat{b} + \psi\hat{g}\right).$$

Since $\tau, \tau', \hat{b}$, and $\hat{g}$ are the same for every agent (the tax schedule can only depend on observable savings, not on unobservable income), the right-hand side is a constant while the left-hand side is strictly monotonic in e whenever $\tau'(\hat{b}) > -1$, because u' is strictly decreasing and $e - \hat{b} - \tau(\hat{b})$ is strictly increasing in e; this holds regardless of whether $\tau'(\hat{b})$ is positive, negative, or zero. In the remaining case, $\tau'(\hat{b}) = -1$, the left-hand side is identically zero

12 Let $f \in \mathcal{C}_2([a, b])$ and define $|| \cdot || : \mathcal{C}_2([a, b]) \to \mathbb{R}_+$ as:

$$||f|| = \max_{x \in [a,b]} |f(x)| + \max_{x \in [a,b]} |f'(x)| + \max_{x \in [a,b]} |f''(x)|,$$

then this is a norm, called the $\mathcal{C}_2$-norm.

13 This does *not* require the objective to be jointly concave in the agent's two choice variables (b, b^c) – a property $\tau \in \mathcal{T}$ does not in general guarantee, since $\tau''(b) < 0$ is admissible. Appendix B (proof of Lemma 2, part 1) instead establishes uniqueness and sufficiency of the first-order conditions directly, by comparing utility across the finitely many candidate solutions consistent with the Karush–Kuhn–Tucker conditions.

14 Section 5.4 argues that non-continuous taxes are also not optimal.

and so cannot equal the strictly positive right-hand side for any e. Either way, the first-order condition can hold for at most one income type, so $\hat{b}$ cannot be an interior optimum for every agent simultaneously: since E contains at least two distinct income types, at least one agent has an incentive to deviate from $\hat{b}$. On top of this, there is an additional complication for the government: now the couch represents a way of savings that are not taxed. Hence, if taxes are sufficiently high, agents may find it optimal to change their bank savings and use the couch as a way of alleviating their tax burden. Hence, taxing savings is in general different than taxing income, whenever the government has no way of directly monitoring income (which is the case in economies with large informal sectors). The remainder of this section gives details on how taxes should be constructed, in order to take care of the incentive-compatibility issues that arise when only savings are observable.

5.1 Equilibrium and Government's Problem

If I denote $\left(b_{\tau,g}(e), b^c_{\tau,g}(e)\right)$ to be the optimal savings for an agent of income type e given τ, g, then the provision of public good is given by:

$$\tilde{g} = \sum_{i=1}^{n} \tau\left(b_{\tau,g}(e_i)\right)\pi_i.$$

One requirement that I impose is for the economy to be in equilibrium given τ, i.e., for there to be a consistency between the public good provision that agents expect, g, and the realized public good provision $\tilde{g}$. In other words, the government must manage a balanced budget. This is a genuinely different problem from the standard optimal taxation setting, where the government's budget requirement is typically an exogenous revenue target evaluated only at the policy ultimately chosen. Here, because g enters agents' old-age consumption directly and therefore shapes their savings decisions, it is not merely a residual that clears once an optimal τ has been identified: it is an equilibrium object that must be solved for jointly with agents' behavior under *every* tax schedule $\tau \in \mathcal{T}$, not only the eventual optimum. This is because comparing the welfare delivered by two candidate schedules requires knowing the equilibrium g that each one induces, so a fixed point must be shown to exist (and be well-behaved) for the entire admissible set $\mathcal{T}$ before the government's optimization problem is even well-posed, rather than being verified once at a single candidate solution. Towards this characterization, I consider the following definition.

Definition 3. Let τ be fixed. A (Nash) equilibrium is $\left(g, \left\{\left(b_{\tau,g}(e), b^c_{\tau,g}(e)\right)\right\}_{e\in E}\right)$ such that:

1. Given τ, g; $\left(b_{\tau,g}(e), b^c_{\tau,g}(e)\right)$ solve:

$$\max_{b,b^c\geq 0} \; u\left(e - b - b^c - \tau(b)\right) + \beta u\left(Rb + b^c + \psi g\right),$$

   for every $e \in E$.

2. There is budget balance:

$$g = \sum_{i=1}^{n} \tau\left(b_{\tau,g}(e_i)\right)\pi_i.$$

Let $\mathcal{G}(\tau)$ be the set of all possible equilibrium outcomes g given τ.

Appendix B shows that if $\tau \in \mathcal{T}$ then $\mathcal{G}(\tau) \neq \emptyset$. Furthermore, the mapping $\mathcal{G} : \mathcal{T} \to \mathbb{R}_+$ is actually a function (i.e., each tax function is assigned exactly to one g) and it is continuous.[15]

With this in mind, the problem that the government solves is to find a tax schedule $\tau \in \mathcal{T}$ such that it maximizes the economy's utilitarian welfare, incentive-compatibility holds for the savings decisions of each income type, and there is consistency between the expected and realized public good provision:

$$\max_{\tau \in \mathcal{T}} \sum_{i=1}^{n} \left[u\left(e_i - \tau(b_{\tau,g}(e_i)) - b_{\tau,g}(e_i) - b^c_{\tau,g}(e_i)\right) + \beta u\left(Rb_{\tau,g}(e_i) + b^c_{\tau,g}(e_i) + \psi g\right)\right] \pi_i, \quad \text{s.t.:}$$

$$\left(b_{\tau,g}(e_i), b^c_{\tau,g}(e_i)\right) \in \operatorname{argmax}_{b,b^c \geq 0}\ u\left(e_i - b - b^c - \tau(b)\right) + \beta u\left(Rb + b^c + \psi g\right) \text{ for every } e_i \in E,$$

$$g = \mathcal{G}(\tau).$$

This program is well defined for every $\tau \in \mathcal{T}$ since $\mathcal{G}(\tau) \neq \emptyset$ and an optimal solution for each agent's problem exists. Furthermore, the assumptions imposed on $\mathcal{T}$ guarantee that this is a compact set within the space of continuous functions,[16] and since the objective function is continuous in τ,[17] then a maximum is attained within the set $\mathcal{T}$.

This problem can be thought of as a principal-agent problem in which the principal (the government) wants to choose a rule (tax schedule) to induce the agents to choose an incentive-compatible action. Hence, a similar problem as the one described by Rogerson (1985) arises with the incentive-compatibility constraints: that they represent a continuum of constraints, which in general may be difficult to handle. Fortunately, since τ is restricted to $\mathcal{T}$, this imposes enough structure on each agent's problem that I can use the first-order approach to characterize the optimal tax schedule. And, as discussed by Rogerson (1985), due to the sufficiency of the first-order conditions, this guarantees that the first-order approach is valid in the context of my model.

5.2 Savings and Consumption Decisions Given Taxes

The main difference between this and the benchmark framework is that the curvature of the tax schedule is important to determine the optimal amount of bank savings for each individual. Taxing too high may push agents out of the formal financial system towards the couch. The following lemma characterizes when an individual chooses to save in the bank or couch as a function of the tax schedule.

15 This function is continuous with respect to the $\mathcal{C}_2$-norm defined in the twice-continuously differentiable functions. For more details consult Appendix B.

16 Since every $\tau \in \mathcal{T}$ is bounded and has a uniformly bounded first derivative (i.e., there exists $M > 0$ such that $|\tau'(b)| \leq M$ for all $\tau \in \mathcal{T}$ and $b \in [0, e_n]$), the family $\mathcal{T}$ is uniformly bounded and, by the Mean Value Theorem, equicontinuous: $|\tau(x) - \tau(y)| \leq M|x-y|$ for every $\tau \in \mathcal{T}$ and $x, y \in [0, e_n]$. Since $[0, e_n]$ is compact, the Arzelà–Ascoli Theorem then implies that every sequence in $\mathcal{T}$ has a uniformly convergent subsequence, i.e., that $\mathcal{T}$ is (sequentially) compact in the space of continuous functions on $[0, e_n]$.

17 The continuity of the objective function as a function of τ is also a consequence of all the differentiability and strict concavity structure imposed in this problem. More details can be found in Appendix B.

Lemma 2. *Let $\tau \in \mathcal{T}$, $g > 0$, and $\left(b_{\tau,g}(e), b^c_{\tau,g}(e)\right)$ is a solution of type $e \in E$'s problem. Then:*

1. *The solution to $e \in E$'s problem is unique, and the first-order conditions are necessary and sufficient to characterize it.*
2. *If $b_{\tau,g}(e) > 0$ and $b^c_{\tau,g}(e) = 0$ then:*
$$\tau'\left(b_{\tau,g}(e)\right) \leq R - 1.$$
3. *If $b_{\tau,g}(e) = 0$ and $b^c_{\tau,g}(e) > 0$ then:*
$$\tau'(0) \geq R - 1.$$
4. *If $b_{\tau,g}(e) > 0$ and $b^c_{\tau,g}(e) > 0$ then:*
$$\tau'\left(b_{\tau,g}(e)\right) = R - 1.$$

The proof of this lemma can be found in Appendix B. This lemma states that the marginal tax rate τ' and its comparison to the interest rate $R - 1$, which is the net return on bank savings, determines whether an individual chooses the couch or the bank to save for the future. An individual chooses the couch whenever the government "taxes away" the benefits of saving in the bank, i.e., whenever the marginal tax rate exceeds the return the bank guarantees.

The precise level of savings in either system is a function of not only τ but also of the agent's income type e and g. The following proposition describes some qualitative features of the optimal savings decision.

Lemma 3. *Let $\tau \in \mathcal{T}$, $g > 0$, and $\left(b_{\tau,g}(e), b^c_{\tau,g}(e)\right)$ is the solution of type $e \in E$'s problem. Then:*

1. *$b_{\tau,g}(\cdot)$ is a non-increasing function of g, while it is a non-decreasing function of e.*
2. *$b^c_{\tau,g}(\cdot)$ is a non-increasing function of g.*

Appendix B provides the proof of this lemma. This lemma states that bank savings cannot decrease with income, suggesting that richer individuals save at least as much as poorer agents in the bank. This is a result of each agent wanting to smooth as much as possible consumption over time (since $\beta R = 1$) and, hence, richer individuals want to transfer more resources to the future to smooth their consumption. The lemma does not establish an analogous monotonicity for couch savings: as the model's low-income-couch/high-income-bank sorting illustrates, couch savings can in fact fall to zero as income rises, once an agent's income is high enough that the bank alone dominates. Importantly, the proof of this result heavily relies on $\tau \in \mathcal{T}$, which guarantees that the decision of b, b^c given τ is smooth and well-behaved in a way that makes bank savings differentiable with respect to income.

In addition, both couch and bank savings are non-increasing in g. When an individual expects a higher provision of the public good, it understands that its consumption when old will increase, and actually, since $\psi > R$, consumption when old will increase more than if either bank or couch savings were to rise. Hence, it is optimal for the agent to reduce (or not increase) their savings, allowing the agent to increase consumption when young and smooth consumption over time.

The following lemma characterizes how optimal consumption choices when young and old depend on the curvature of the tax schedule.

Lemma 4. *Let $\tau \in \mathcal{T}$, $g > 0$, and $\left(b_{\tau,g}(e), b^c_{\tau,g}(e)\right)$ be the solution of income type e's problem. Let $c^y_{\tau,g}(e) = e - \tau(b_{\tau,g}(e)) - b_{\tau,g}(e) - b^c_{\tau,g}(e)$ and $c^o_{\tau,g}(e) = Rb_{\tau,g}(e) + b^c_{\tau,g}(e) + \psi g$ be the optimal consumption when young and old (respectively) of agent $e \in E$.*

1. *If $b_{\tau,g}(e) > 0$ and $b^c_{\tau,g}(e) = 0$ then:*
   (a) *$c^y_{\tau,g}(e) > c^o_{\tau,g}(e)$ if and only if $\tau'(b_{\tau,g}(e)) > 0$.*
   (b) *$c^y_{\tau,g}(e) < c^o_{\tau,g}(e)$ if and only if $\tau'(b_{\tau,g}(e)) < 0$.*
   (c) *$c^y_{\tau,g}(e) = c^o_{\tau,g}(e)$ if and only if $\tau'(b_{\tau,g}(e)) = 0$.*
2. *If $b^c_{\tau,g}(e) > 0$ then $c^y_{\tau,g}(e) > c^o_{\tau,g}(e)$.*

The first part of this lemma states that if an agent decides to save b, and the tax schedule is increasing at b, then the agent must consume more when young than when old. For an agent to decide to consume the same amount in both periods, the tax rate has to be constant at b, which means that the distortion introduced by taxes is not relevant for the agent (as in the benchmark model). Finally, the last part of the lemma implies that if an agent has positive couch savings, it must be the case that she is consuming more when young. The proof of this can be found in Appendix B.

5.3 Optimal Tax Schedule

Taxes reduce the attractiveness of the formal financial system by diminishing the benefits of earning a higher return through bank savings, as part of these returns are appropriated by the government. From a utilitarian welfare perspective, however, diverting agents away from bank savings entirely, towards informal saving methods like storing money under the couch, eliminates any gain from higher interest rates and reduces old-age consumption. Moreover, when agents avoid the formal system, it becomes more difficult for the government to observe their incomes, limiting its capacity to levy taxes. This, in turn, lowers public good provision and further depresses old-age consumption. Consequently, a welfare-maximizing government should implement a tax schedule that promotes bank savings and discourages evasion. This section characterizes how such a policy can be optimally designed.

The first result towards characterizing the optimal tax schedule is presented in Proposition 2. This result states that any tax schedule $\tau \in \mathcal{T}$ that induces at least one agent to save on both the couch and the bank (i.e. being an evader) cannot be optimal.

Proposition 2. *Let $\tau \in \mathcal{T}$ be such that there is at least one $e \in E$ with $b_{\tau,g}(e) > 0$, $b^c_{\tau,g}(e) > 0$, and $g = \mathcal{G}(\tau)$. Then τ cannot be optimal.*

In a nutshell, this proposition states that an optimal tax implies zero evasion (evasion is understood as an agent saving in both the couch and the bank). Intuitively, if a tax schedule τ induces agent e to have positive savings in both the couch and the bank, it is because the marginal tax rate at the optimal b is equal to $R-1$, which means that the government is "taxing away" all the benefits of bank savings. Hence, if the government considers a different tax schedule $\hat{\tau}$ that marginally decreases taxes and $\hat{\tau}'(b) < R-1$, then the agent will find it optimal to substitute bank for couch savings, increasing agent e's welfare. Now, for this argument to be complete, I need to show that this change to $\hat{\tau}$ does not incentivize other agents to change their savings structure drastically, and also that the contribution of public good will not decrease when the tax schedule is $\hat{\tau}$. Appendix B shows that if one constructs $\hat{\tau}$ carefully, it is possible to find a feasible tax schedule that will increase at least type e's welfare, without decreasing the provision of the public good. This represents a strict increase in total welfare, implying that τ was not optimal.

The previous proposition suggests that it is optimal for the government to induce agents to use either the bank or the couch but not both systems simultaneously. As I discuss in the following section of the paper, depending on the model's parameter values, the government may find it optimal to allow the lowest-income individuals to use only the couch (which I interpret as a form of financial informality). Whether low-income individuals become financially informal under the optimal tax schedule depends heavily on how much revenue the government can extract from higher-income agents. If the economy has a left-skewed income distribution (i.e., a high mass of agents are richer), then the government can allow itself to exclude the poorest agent from the formal financial system. On the other hand, economies with many individuals at the bottom of the income distribution will have to rely on taxing everyone, since this enables the provision of public goods.

The next thing to emphasize is that this problem has multiple solutions. The multiplicity arises from how taxes may behave outside of the agents' bank savings equilibrium choices. In particular, let $\bar{b}$ be the largest bank savings choice induced by a tax schedule τ and a given $g \in \mathcal{G}(\tau)$. By "revealed preference," none of the agents choose $b > \bar{b}$, since either taxes are too high for those savings levels or they do not need to save as much given the level of public good provision. Consider a new tax schedule $\tilde{\tau}$ such that $\tilde{\tau}(b) = \tau(b)$ for all $b \leq \bar{b}$ and $\tilde{\tau}(b) > \tau(b)$ for $b > \bar{b}$. Let agents take as given $g \in \mathcal{G}(\tau)$. Under the new tax schedule, agents choose exactly the same bank savings that under τ, since $\tilde{\tau}$ is charging even higher taxes in the region that no agent was choosing under τ. As a consequence, $\mathcal{G}(\tilde{\tau}) = \mathcal{G}(\tau)$ and agents have the same bank and couch savings choices. In particular, this means that even if τ is an optimal tax schedule, one can always construct a new tax schedule that charges higher taxes "at the right tail" and is also optimal. This construction, which Figure II alludes to, in which the blue curve may be an optimal tax, but one can always construct a new tax schedule (the green) which is also optimal.

Before I discuss the main theoretical result of the paper, I present two examples that provide an insight into why the optimal tax schedule cannot be strictly decreasing or strictly increasing and convex. Figure III displays the optimal response of the agents across the income distribution to a tax schedule that is decreasing. Since $\tau'(\cdot) < 0$ in its domain, then $\tau'(b) \leq R-1$ for every $b \geq 0$, which implies that all agents in the economy will have zero savings in the couch. Since the tax rate is decreasing, this incentivizes richer individuals to save more in the bank, which from a social point of view is what the government would want. However, this tax structure is very regressive in the following sense: most tax revenue is collected from the poorest individuals, and richer agents are contributing less to tax revenue (they are saving more but

Figure II: Optimal tax schedule is not unique.

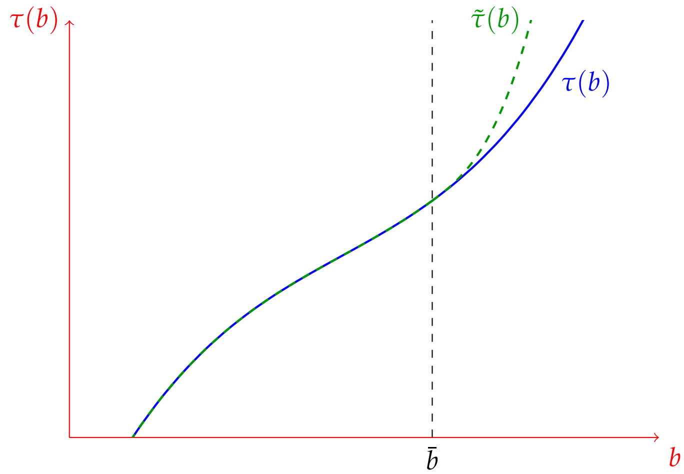

taxes are decreasing in savings). The poorest agents are being so heavily taxed that in this example, they even decide not to save anything in the bank since their consumption when young is very low. The provision of the public good is a way for the government to redistribute part of the income in the economy, and a decreasing tax schedule is creating more inequality across different individuals. The poorest could do better if the government allows them to consume more when young, which can be achieved by reducing taxes.

Figure III: A Decreasing Tax Schedule Is Regressive.

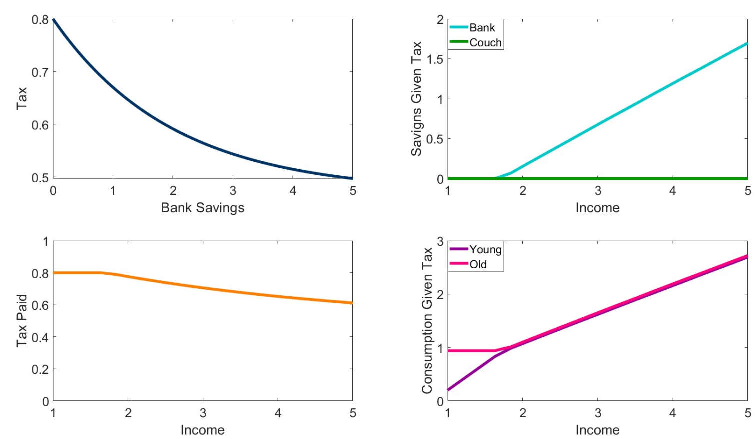

Notes: The top-left panel presents the tax schedule, which is decreasing and convex. The top-right panel presents the savings decisions for the bank (light blue) and the couch (green) induced by the tax schedule. The bottom-left panel presents the tax paid by each income type, while the bottom-right panel presents consumption when young and old. For this figure, I consider $R = 1.05$, $\psi = 1.25$, and $E \sim U([1,5])$.

All decreasing tax structures (regardless of their curvature) will induce a similar response: richer agents will pay fewer taxes and save more money in the bank. This is why the optimal tax schedule is not a decreasing function of bank savings (at least intuitively; a formal argument is provided below).

Figure IV presents the equilibrium response given a progressive τ. Progressive taxes must be convex, and this figure presents the case of a strictly increasing and convex tax schedule. This tax schedule has the property that $\tau'(\tilde{b}) = R - 1$ with $\tilde{b} = 1$, and hence by Lemma 2 every agent will choose to save an amount lower or equal than $\tilde{b}$. Since savings are non-decreasing with e and this tax schedule is increasing, in theory richer agents should contribute more to the provision of the public good. However, it is not optimal for agents to save in the bank more than $\tilde{b}$, which implies that all rich agents will pool at $b = \tilde{b}$ and save the rest in the couch, as can be seen in the top-right panel of this figure. Hence, this tax schedule creates evasion at the top of the income distribution. And, as stated in Proposition 2, if τ induces evasion, then it cannot be optimal.

Figure IV: A PROGRESSIVE TAX SCHEDULE INDUCES EVASION.

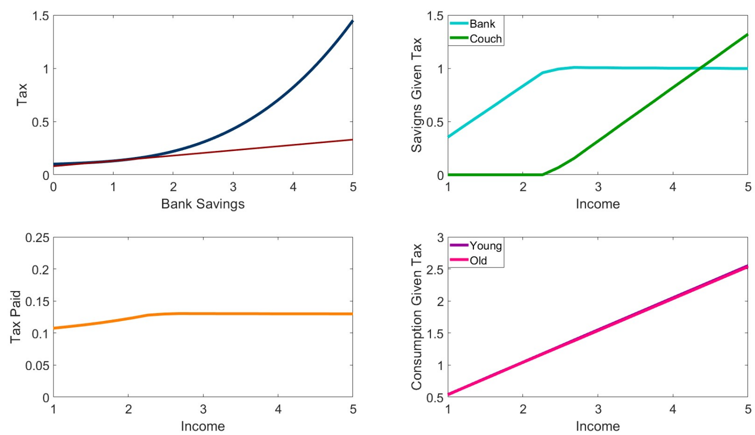

NOTES: The top-left panel presents the tax schedule, which is progressive (increasing and convex). The red line in this panel represents the tangent at $b = 1$, with a slope equal to $R - 1$. The top-right panel presents the savings decisions for the bank (light blue) and the couch (green) induced by the tax schedule. The bottom-left panel presents the tax paid by each income type, while the bottom-right panel presents consumption when young and old. For this figure, I consider $R = 1.05$, $\psi = 1.25$, and $E \sim U([1,5])$.

Progressive taxation disincentivizes richer agents away from high levels of savings in the bank, which, from a social welfare point of view, is not desirable since many resources are not being allocated towards the more efficient way of savings. A similar behavior will be induced by any progressive tax schedule that is redistributive enough to reach the point at which the marginal tax rate absorbs the full return on formal saving, which is why the optimal schedule within the concave class I characterize is not progressive (a formal argument is given below). This is in direct contrast to Vairo (2025), which features asymmetric information and a government that cares about both utilitarian welfare and tax revenue. In Vairo (2025), the government is uncertain about agents' income-generating processes and seeks to maximize the average of utilitarian welfare and tax revenue, considering the worst-case scenario for them. The main result is that the (robustly) optimal tax schedule is convex, since it aligns both the private sector's interest (convex taxes imply a concave consumption function and hence make them risk averse) and the government's preferences for tax revenue, which are also risk averse. My model is one in which there is also hidden information for the government, however in my framework agents have an "outside option" which can be used (although with the cost of losing R), and hence the government must internalize this fact given that it cares about both welfare but also tax revenue (that materializes in the provision of a public good).

In what follow of the theoretical characterization of optimal taxes, I restrict attention to concave tax schedules. Concavity is a natural restriction to impose on the space of admissible tax schedules, in the same spirit as the other restrictions already embedded in $\mathcal{T}$: it guarantees a non-increasing marginal tax rate, and it delivers the tractable characterization of the optimal schedule in Proposition 4 below.

Proposition 3. *There exists a tax schedule $\tau^\star \in \mathcal{T}$ that is concave and maximizes utilitarian welfare among concave tax schedules.*

Note that a non-increasing marginal tax rate is the opposite curvature from the progressivity notion of Definition 1, which requires a non-decreasing marginal tax rate: a concave τ is not progressive in that sense, even though its tax liability can still be increasing in b, so richer agents can still pay weakly more in absolute (though not marginal) terms. I do not prove that restricting attention to concave schedules is without loss of generality relative to the unconstrained problem over all of $\mathcal{T}$, so the characterization below should be read as identifying the optimal tax schedule *within the concave class*, not as an unconditional optimum among all admissible schedules. As a numerical check, however, I solved the government's problem without imposing concavity, using the same Chebyshev-polynomial approximation described in Appendix C; across every parameterization I considered, the resulting optimal schedule was concave, which is suggestive evidence that the restriction is not binding in practice, even though I do not establish this analytically.

The proof of Proposition 3 establishes existence directly: the set of concave schedules in $\mathcal{T}$ is compact once endowed with the topology of uniform convergence of τ and τ', and welfare is continuous in this topology, so a maximizer exists by the extreme value theorem. Appendix B provides the details.

I now present the main result of the paper, which is a proposition that characterizes an optimal concave tax schedule.

Proposition 4. *Let $\tau^\star$ be an optimal concave tax schedule. Then:*

1. *There exists $\tilde{b}$ such that $\tau^\star(b) = \bar{\tau}$ for all $b \geq \tilde{b}$.*
2. *$\tau^\star$ is non-decreasing.*
3. *At $\tau^\star$ there is zero evasion.*
4. *Every agent whose optimal savings are greater than or equal to $\tilde{b}$ consumes the same amount when young and old. However, this amount cannot be equal across all agents with savings bigger or equal to $\tilde{b}$.*

The fact that $\tau^\star$ becomes constant for large savings is due to the government wanting to incentivize richer agents to only save in the bank (since $\tau'(b) = 0 < R - 1$). Now, since $\beta R = 1$ then the utility-maximizing thing for agents to do is to perfectly smooth consumption over time (i.e., consuming the same amount in both periods). If $\tau'(b) \neq 0$, then as Lemma 4 states, this introduces a distortion in the agent's consumption decisions, resulting in consumption varying across time. Hence, to eliminate this distortion, $\tau'(b) = 0$ at the top, which guarantees that richer agents will be able to consume the same amount when young and old.

This result is similar to the one discussed in Mirrlees (1971), where the author shows that, under suitable conditions, optimal income taxes with private information about worker productivity have a zero marginal tax rate for the highest type. Intuitively, this happens because there is no higher type above the top whose incentive to mimic needs to be deterred, so a marginal tax cut for the highest type is a costless welfare improvement. Although similar, my result stems from a very different force: in my model, agents have an "outside option" that becomes more attractive as taxes increase. Hence, the government's ability to tax is bounded by the gain between the bank and outside option $R - 1$. Furthermore, the government wants people to contribute to the provision of the public good since this is a mechanism to redistribute income in the economy, and the way of maximizing contribution is to eradicate evasion, which can be done if taxes eventually become constant.

Figure V: Optimal Concave Tax Schedule $\tau^\star(b)$.

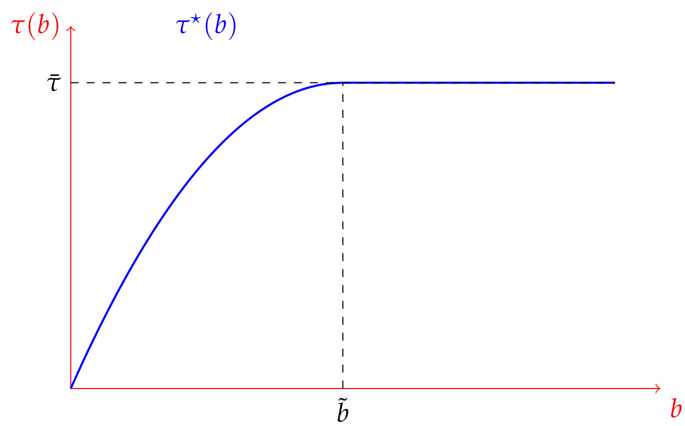

By Lemma 3, bank savings are already non-decreasing in income under *any* admissible $\tau \in \mathcal{T}$; the optimal tax's own non-decreasingness (established in the proof of Proposition 4 via concavity together with flatness at the top) is a separate feature of $\tau^\star$, not the cause of the savings monotonicity. Together, the two facts imply that richer agents weakly save more in the bank and pay weakly more in taxes, which is a way of redistributing income across agents; since Proposition 4 pools every agent at or above $\tilde{b}$ on the flat tail, this monotonicity is generally weak rather than strict, and savings and tax payments can be exactly equal across a range of income types. In general, the government wants more agents to be saving in the bank since it has a higher return than the couch. For an agent to be willing to choose the bank over the couch, the marginal tax rate has to be below $R - 1$. Since the marginal tax rate is non-increasing for a concave tax schedule, if one agent chooses to save in the bank, then every agent with a higher income will also choose to save in the bank. Also, as b approaches zero, this increases the marginal tax rate, potentially above the interest rate, which would invite very low-income agents to save in the couch. As Lemma 4 states, an agent who is choosing to save in the couch is necessarily consuming more when young than when old; this tilt is exactly what the agent's own first-order condition, $u'(c_y) = \beta u'(c_o)$, already equates at the chosen allocation, so it is not itself an additional source of welfare gain from $\beta < 1$. Whether inducing an agent to save in the couch raises welfare instead depends on the full comparison across tax schedules, including the higher return to formal saving, tax burdens, revenue, and public-good provision. Indeed, as Section 6 highlights, the government pushing some agents towards the couch may not always be optimal, particularly if the income distribution is right-skewed (i.e., the mass of low-income agents is elevated).

This zero marginal tax rate at the top property of the optimal tax implies that every agent that saves $b \geq \tilde{b}$ will consume the same amount when young and old. However, this amount cannot be equal across all such agents. Otherwise, the marginal tax rate would be positive, and the tax structure would be progressive (as in the full information case).

5.4 Step Tax Schedules Are Not Optimal

The concave structure of the optimal tax provides agents with an incentive to distinguish between financial informality and bank savings. As discussed in the quantitative section (next part of the paper), it may even be optimal for the government to charge zero taxes to financially informal agents. This section argues that the concave shape between $\tau^\star(0)$ and $\bar{\tau}$ is necessary to keep incentives. One could think that a simpler tax schedule, such as a tax that is a step function, that is zero until $\tilde{b}$ and then equal to $\tilde{\tau}$ may be able to implement the same allocation as $\tau^\star$.[18] But, as Figure VI suggests, this is not the case. The main problem with this tax schedule is that it creates evasion: agents who under $\tau^\star$ were saving b that was above but close to $\tilde{b}$ are now going to be incentivized to save $\hat{b} < \tilde{b}$, and then by $\tau^\star(0)$ taxes (which in this example is equal to zero), and then save the same level as before, but saving $b - \hat{b}$ in the couch.

As a result of agents evading, the total provision of public goods is going to be lower under this new tax schedule than under $\tau^\star$, which decreases the utility of all agents in the economy, but particularly harm low income individuals, who were heavily benefited from the provision of the public good. In the example presented in Figure VI, the aggregate tax collection is zero since all agents find it optimal to save now $b < \tilde{b}$.

Furthermore, if I look for the "best step function", in the sense of looking for the step tax schedule that maximizes utilitarian welfare, welfare under the best step schedule appears to rise with the number of allowed steps, and the resulting schedule appears to approach the optimal tax schedule characterized by Proposition 4. Figure VII illustrates this for the richest step-function specification I solved for, allowing five steps: the optimal five-step schedule (aqua) closely tracks the concave characterization (blue).

This exercise suggests that restricting attention to tax schedules in $\mathcal{T}$, which contains continuous and differentiable functions, is without loss of generality. While this restriction is imposed to ensure the validity of the first-order approach, the example indicates that even when optimizing over functions outside of $\mathcal{T}$, the resulting solutions tend to approximate those within

18 This tax function is not continuous nor differentiable, and hence it is not an element of $\mathcal{T}$. However, since taxes are often step functions in many economies, I consider this class of taxes important and discuss why they are not optimal in the presence of hidden savings.

Figure VI: STEP-TAXES INDUCE EVASION.

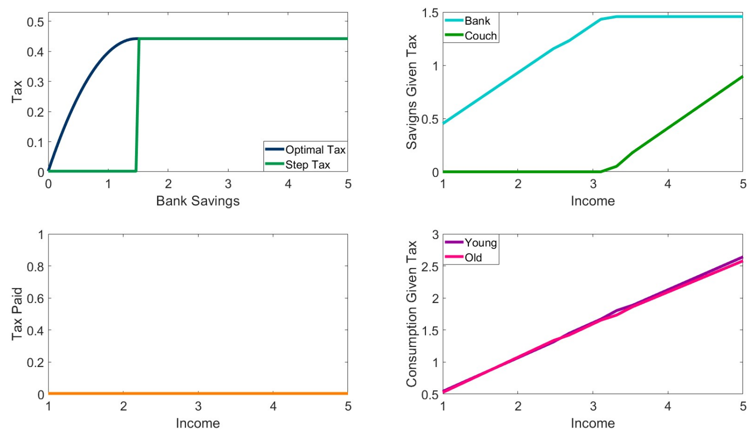

NOTES: The top-left panel presents the optimal tax schedule (in blue) and an example of a step tax (in aqua). The top-right panel presents the savings decisions for the bank (light blue) and the couch (green) induced by the tax schedule. The bottom-left panel presents the tax paid by each income type, while the bottom-right panel presents consumption when young and old. For this figure, I consider $R = 1.15$, $\psi = 1.25$, and $E \sim \text{BetaBin}(20, 3, 1)$. This density is described in Appendix C.

Figure VII: STEP-TAXES APPROXIMATE THE OPTIMAL CONCAVE TAX.

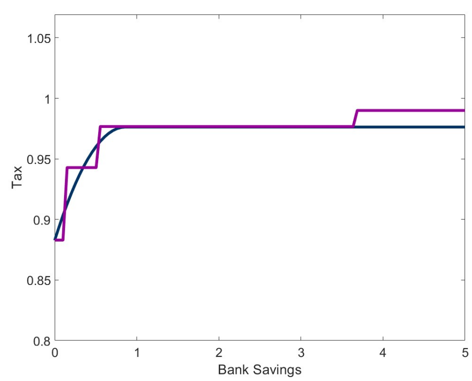

NOTES: The top-left panel presents the optimal concave tax schedule implied by Proposition 4 (in blue) and the optimal step tax when five steps are allowed (in aqua). For this figure, I consider $R = 1.15$, $\psi = 1.25$, and $E \sim \text{BetaBin}(20, 3, 1)$. This density is described in Appendix C.

the set.

6 Quantitative Analysis

This section aims to understand how the model's primitives (in particular, the income distribution and the interest rate) shape the optimal tax schedule and the induced response from agents. To do this, I use the numerical solution of the government's problem since, in general, it does not have a closed-form solution. The government's problem is a variational problem in which it is choosing a tax function to maximize utilitarian welfare. Hence, I use Ritz's method to solve this numerically and discretize the problem. Details can be found in Appendix C.

Given the model's parameters, the model not only has implications for the optimal tax policy, but it also allows me to analyze (by analyzing each agent's response to the tax schedule) the savings' distribution (for both the bank and couch), consumption distribution, and the financial informality rate. If an agent has positive couch savings but zero bank savings, I interpret this agent as financially informal. On the other hand, if an agent has positive bank and couch savings, I interpret this as evasion. By Proposition 2, the optimal tax schedule implies zero evasion; however, as I show in this section, the informality rate may not be zero.

Figure VIII presents the optimal tax schedule as well as the induced behavior for the case of a uniform income distribution. The government wants to be able to collect as much taxes as possible while also pushing as many agents as possible towards the formal sector (bank) and allow them to smooth consumption over time. To do this, the government must charge a constant tax $\bar{\tau}$ if savings are greater than $\tilde{b}$. Now, charging a flat tax equal to $\bar{\tau}$ over all possible savings levels is not optimal as well, since this induces low income agents to actually not save in either system ($\bar{\tau}$ is too high of a tax for them to pay). What is actually optimal is to have $\tau^{\star}(0) < \bar{\tau}$ and concavely increase taxes until the level $\bar{\tau}$. I want to highlight two things here: first, since the marginal tax rate is larger than $R-1$ at $b = 0$ ($\tau^{\star'}(0) > R-1$) then low income agents find it optimal to save in the couch, which allows them to transfer some resources between periods while still consuming more when young (every agent that uses the couch always has higher consumption when young). Second, the fact that $\tau^{\star}$ concavely increases up to $\bar{\tau}$ ensures that agents that have bank savings slightly higher than $\tilde{b}$ actually want to remain saving in the bank and not become evaders.[19]

As a result of this tax schedule, some agents exclusively save on the couch, and others only save on the bank. Hence, given the income distribution, even at the optimal tax schedule, the informality rate is positive (it is around 30%). This also has as a result a non-decreasing effective tax function, as seen in the bottom-left panel of Figure VIII, in the sense that richer agents do not pay fewer taxes than low-income agents.

I now analyze the shape of the optimal tax schedule when the income distribution is left-skewed, i.e., a large mass of agents have high income. The optimal tax and induced behavior are presented in Figure IX. This economy has a higher average income than the one with a uniform distribution. Hence, the government can allow itself to be more "progressive" when designing the optimal tax, in the sense that it charges lower taxes (actually zero taxes) to agents who save zero in the bank and uses the tax collected from richer agents to finance the public good provision. In the uniform case, the government could not charge zero to low-income agents since there is a significant mass of low-income agents, and charging zero taxes would dramatically decrease public good provision. However, in the left-skewed case, the government can allow itself not to collect taxes from low-income agents. Again, the shape of the tax schedule for $0 < b < \tilde{b}$ is such that it incentivizes richer agents to save in the bank and not mimic low-income types and pay zero taxes.

Even with this income distribution, the optimal informality level is not zero. In this case, the informality rate is around 3%. Hence, economies with a larger mass of high-income individ-

19 Section 5.4 discusses this in more detail.

Figure VIII: Optimal Tax Schedule Under Uniform Distribution.

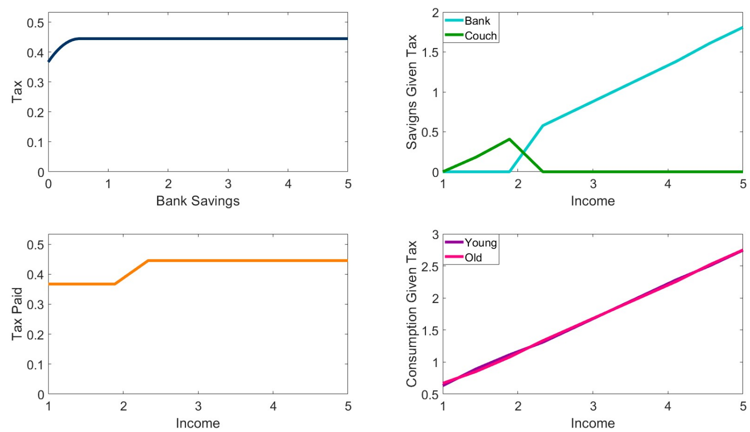

NOTES: The top-left panel presents the optimal tax schedule. The top-right panel presents the savings decisions for the bank (light blue) and the couch (green) induced by the tax schedule. The bottom-left panel presents the tax paid by each income type, while the bottom-right panel presents consumption when young and old. For this figure, I consider $R = 1.15$, $\psi = 1.25$, and $E \sim U([1,5])$.

Figure IX: Optimal Tax Schedule Under Left-Skewed Distribution.

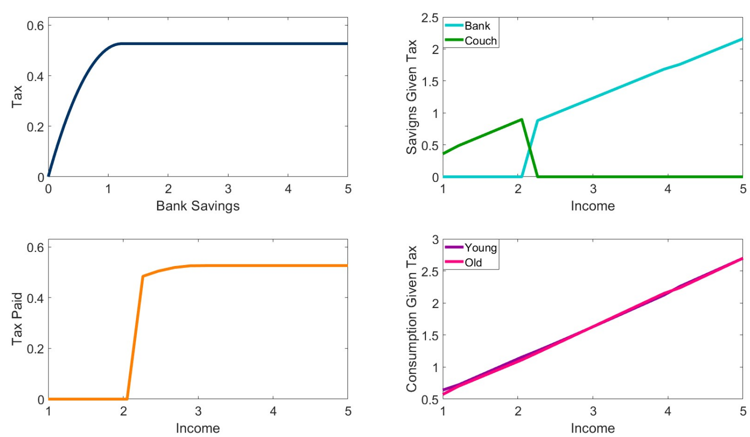

NOTES: The top-left panel presents the optimal tax schedule. The top-right panel presents the savings decisions for the bank (light blue) and the couch (green) induced by the tax schedule. The bottom-left panel presents the tax paid by each income type, while the bottom-right panel presents consumption when young and old. For this figure, I consider $R = 1.15$, $\psi = 1.25$, and $E \sim \text{BetaBin}(20, 3, 1)$. This density is described in Appendix C.

uals can rely on collecting taxes from wealthier individuals while sparing very low-income individuals from taxation.

Table I summarizes the behavior induced by the optimal tax schedule under different parametric assumptions of the income distribution and the interest rate. All of these exercises are considering a Beta-Binomial distribution, whose description can be found in Appendix C.

Table I: PUBLIC GOOD PROVISION AND FINANCIAL INFORMALITY ACROSS DIFFERENT ECONOMIES.

Income Distribution	Description	Interest Rate	$g^\star$	Fin. Informality Rate
$\alpha = \beta = 1$	Uniform	$R = 1.15$	0.42	30%
$\alpha = 1, \beta = 3$	Right-Skewed	$R = 1.15$	0.34	66.92%
$\alpha = 3, \beta = 1$	Left-Skewed	$R = 1.15$	0.49	3.64%
$\alpha = \beta = 1$	Uniform	$R = 1.05$	0.15	50%
$\alpha = 1, \beta = 3$	Right-Skewed	$R = 1.05$	0.12	84.621%
$\alpha = 3, \beta = 1$	Left-Skewed	$R = 1.05$	0.14	12.59%

The main takeaway of this table is that, for a fixed value of R, as an economy has a higher mass of high-income agents, the government will be able to collect more taxes, provide higher public goods, and induce a lower informality rate by using the optimal tax schedule. On the other hand, if we fix the income distribution, an economy with lower interest rates can collect fewer taxes and have higher informality rates. This is because $R - 1$ represents an upper bound on the marginal taxes the government can charge and incentivize agents to save in the bank. If marginal taxes are above $R - 1$, then agents will choose to evade paying lower taxes. Then, if the interest rate decreases, so does the capacity of the government to charge high taxes.

6.1 An External Check: Cross-Country Patterns in Formalization

The exercises above are illustrative: they show how the model's primitives shape the optimal tax schedule and the resulting financial informality rate, but they say nothing about whether the model's central comparative static—that a higher return to formal saving is associated with greater financial formalization—is visible outside the model itself. This section takes a modest, disclosed step in that direction, without calibrating the model to any single economy.

As discussed in Section 3, R in the model is defined relative to the couch, whose return is normalized to one; the theory requires only a positive spread between the bank's return and this informal alternative, not that either return take any particular absolute value. I proxy this spread using each country's average *nominal* return on bank deposits: hiding currency preserves its nominal face value exactly, which is the direct empirical counterpart to the couch's normalized return of one. An inflation-adjusted reading is also possible if one instead interprets the couch as a hard asset that tracks purchasing power rather than literal cash; I adopt the nominal reading here as the more literal match to the model's description of the couch. I proxy financial formalization using Account Ownership, the share of adults with an account at a financial institution or mobile-money provider, from the Global Findex Database (Klapper et al., 2025); higher account ownership corresponds to lower financial informality.

I focus on five economies, chosen for data completeness rather than for how well they fit the pattern: Mexico, given it is normally used as a classic example of an economy with a large informal sector; Brazil, Latin America's largest economy; and one large economy from each of three other developing regions—Kenya (Sub-Saharan Africa), Indonesia (East Asia and the Pacific), and Morocco (the Middle East and North Africa). Table II reports account ownership, the average nominal deposit rate, and the implied value of R for each country.[20] This table is suggestive of the findings that the previous quantitative analysis showed: that as the interest rate R decreases, one should expect to observe a lower financial formalization of the economy.

Table II: NOMINAL RETURNS TO FORMAL SAVING AND ACCOUNT OWNERSHIP.

Country	Account Ownership	Avg. Nominal Deposit Rate	Implied R
Kenya	78.5%	7.58%	1.076
Brazil	74.1%	7.63%	1.076
Indonesia	45.6%	6.31%	1.063
Morocco	44.4%	2.57%	1.026
Mexico	41.4%	2.04%	1.020

NOTES: Account ownership is the share of adults (age 15+) with an account at a financial institution or a mobile-money provider, from the Global Findex Database (Klapper et al., 2025). The average nominal deposit rate is computed using the World Bank's World Development Indicators (World Bank, 2024). Both series are averaged over the same calendar years for each country; see the text for the exact construction. Implied R equals one plus the average nominal deposit rate.

This pattern should not be overstated. With five economies, it is suggestive rather than a formal test. Financial formalization plainly depends on margins the model abstracts from entirely, such as payments infrastructure, regulatory mandates, and macroeconomic stability, and a single-instrument model of the return-to-saving tradeoff was never intended to be a complete account of cross-country variation in financial inclusion. What the five-country comparison suggests is only that the model's central margin is not absent from the data, even though it is evidently incomplete.[21]

6.2 Policy Takeaways

The quantitative exercises highlight several policy-relevant implications that emerge naturally from the structure of the model. A first and central lesson is that *zero financial informality is not generally optimal*, even for a welfare-maximizing government. In economies where income is unobservable and savings can be hidden, financial informality arises as an equilibrium response to taxation rather than as a pure inefficiency to be eliminated. The optimal tax schedule trades off revenue extraction against participation in the formal financial system, and in many parameterizations this trade-off is resolved by allowing a positive mass of low-income agents to remain financially informal. From a normative perspective, this implies that policies aimed at fully eliminating financial informality may be welfare-reducing when they ignore agents' outside options and the distortionary effects of taxing observable financial activity. This lesson concerns financial informality specifically, as modeled here; I do not take a stance on policies

20 Both series are averaged over 2014–2023, using only the specific calendar years in which a Global Findex wave is available for that country (2014, 2017, and 2021, or 2022 for Mexico), so that account ownership and the deposit rate are computed over exactly the same years.

21 A more demanding complementary exercise would calibrate the model to a single country directly—for instance, setting R to Mexico's own return on formal savings and the income distribution to match household income data from Mexico's ENIGH survey—and comparing the model's *optimal*-policy-implied informality rate against an untargeted empirical moment, such as the share of Mexican households reporting saving through an informal savings club or with a person outside the family (Global Findex variable `fin17b`), rather than the account-ownership margin used here. I leave this closer, single-country validation to future work.

targeting labor-market or firm informality, which operate through different margins not captured by this framework.

A second implication is that the *extent of financial informality depends critically on the distribution of income*. The quantitative results show that economies with a larger mass of high-income agents can sustain higher public good provision while tolerating lower financial informality rates, as the government can rely more heavily on taxing richer households without inducing widespread exit from the formal financial system. Conversely, when the income distribution is right-skewed and a large fraction of agents is relatively poor, optimal policy involves taxing a broader base and accepting substantially higher financial informality. This suggests that cross-country differences in the extent of financial informality need not reflect differences in enforcement capacity or institutional quality alone, but may instead be tightly linked to underlying distributional characteristics of income.

More broadly, the model implies that *financial informality should be viewed as an endogenous margin of adjustment*, rather than as evidence of policy failure. In equilibrium, informal saving allows low-income agents to smooth consumption when formal participation would expose them to high marginal tax rates. By tolerating some financial informality at the bottom, the government relaxes participation constraints and avoids inducing inefficient evasion at the top. In this sense, financial informality plays a role analogous to participation constraints in classic screening problems: it restricts the set of implementable allocations and shapes the optimal tax schedule.

The analysis also highlights the importance of *financial returns and financial development* for tax capacity. Since the net return on formal saving places an upper bound on marginal taxation, economies with lower returns to formal saving face tighter constraints on optimal taxation and exhibit higher equilibrium financial informality. Policies that increase the relative attractiveness of formal saving—such as improving financial access or reducing intermediation costs—can therefore expand the feasible set of tax schedules the government can implement, allowing it to sustain higher public good provision without pushing agents out of the formal financial system.

7 Robustness and Generalizations

Throughout the paper, I have relied on several simplifying assumptions to keep the analysis tractable, and it is worth discussing how each of them shapes the results and what would change if they were relaxed.

The normalization $\beta R = 1$ (Assumption 1) is the most pervasive of these, appearing throughout both the benchmark and hidden-savings models, but it is not essential to the paper's qualitative conclusions. Its main purpose is technical: it delivers closed-form expressions for each agent's optimal savings and, in the benchmark model, implies that agents perfectly smooth consumption across periods. If instead $\beta R \neq 1$, the coefficient βR no longer drops out of the agents' first-order conditions, and consumption is no longer exactly equalized across periods even absent taxation. The change, however, is quantitative rather than qualitative: the government still has no reason to distort the intertemporal margin of the richest agents in the hidden-savings model, so the optimal tax schedule continues to become flat ($\tau' = 0$) at the top. What changes is only the interpretation of this flatness. Rather than implying equal consumption when young and old, it now implies that the agent's Euler equation is left undistorted relative to what she would choose absent the tax, preserving her own βR-implied consumption tilt rather than forcing $c^y = c^o$.

A related, and in some ways more consequential, simplification is the assumption $R > 1$ itself. As discussed in Section 3, this is a normalization rather than a substantive restriction: what the model requires throughout is that the bank offer a strictly higher return than the couch, not that either return take any particular absolute value. Since the couch's return is set to one purely as a choice of units, R should be understood as measuring the bank's return *relative to* the informal alternative—a positive spread—rather than as an interest rate in some absolute, economy-wide sense. Every result in the paper depends only on this spread being positive; nothing hinges on R exceeding one in whatever units a reader might otherwise be tempted to interpret it.

This observation also clarifies why the paper treats R as exogenous throughout, rather than as an object determined in general equilibrium. Because R is fundamentally a spread between two returns, closing the model in general equilibrium would require specifying how *both* the bank's return and the couch's return respond to the economy's aggregate savings behavior, which is well beyond what is needed to characterize the shape of the optimal tax schedule analytically. Even setting the couch aside, closing only the bank's side of the model in the most natural way—treating aggregate bank savings as clearing a closed credit market with no outside investment technology—would require that aggregate savings net to zero, since one agent's deposit would have to be borrowed by another for the market to clear. A genuine general-equilibrium closure would therefore require either a real investment technology absorbing aggregate savings—for instance, a production sector whose diminishing returns to capital cause R to fall as aggregate savings rise—or a departure from the no-borrowing assumption altogether, and in either case the fixed point between the optimal tax schedule and the equilibrium return would generally have to be solved numerically rather than characterized analytically. The partial-equilibrium approach adopted here is what allows the paper to obtain sharp, closed-form characterizations of the optimal tax schedule's shape; a full general-equilibrium treatment is a natural, but substantially more demanding, extension.

Finally, the government's commitment to its announced tax schedule, maintained throughout the paper, is not a minor simplification but an assumption that fundamentally determines the model's scope of analysis. Without commitment, once agents have already revealed their savings choices to the government by depositing in the bank, taxing those balances further does not distort a decision that has already been made, so a government free to reoptimize ex post would have an incentive to expropriate observed bank savings well beyond what it would have chosen to announce ex ante. Anticipating this, agents would shift savings toward the couch precisely to avoid being caught by a future, more confiscatory tax, unraveling the very formal-sector participation the optimal tax schedule is designed to sustain. Analyzing this possibility properly would require replacing the mechanism-design approach used here with a dynamic game between the government and agents, along the lines of the sustainable-plans and reputation literatures, which is a substantially different exercise from the one undertaken in this paper.

8 Conclusion

This paper develops a theoretical framework for understanding optimal tax design and public good provision in economies where income is unobservable and savings can be hidden. The model captures the endogenous emergence of financial informality and tax evasion by introducing the possibility of agents choosing between taxed bank savings and untaxed informal savings. The analysis shows that the optimal tax schedule must balance efficiency and redistribution while acknowledging agents' outside options. In particular, a tax schedule that is progressive enough to tax away the full return on formal saving becomes counterproductive, incentivizing wealthier individuals to avoid the formal financial system, limiting revenue and redistribution potential.

Within the class of concave tax schedules, for which I provide a full characterization, the optimal tax structure is flat beyond a threshold, a form that simultaneously deters evasion at the top and limits financial informality at the bottom. I do not establish that concavity is without loss of generality relative to the full space of admissible schedules, though numerical evidence is suggestive that this restriction is not binding. This policy design ensures that high-income individuals continue participating in the formal financial system without being excessively taxed. In contrast, low-income individuals are taxed minimally or excluded from formal taxation to preserve participation. Notably, this approach leads to a segmented equilibrium where some agents remain financially informal, but evasion (saving in both systems) is optimally eliminated. The results highlight how careful tax design can promote formalization, improve welfare, and support public goods even in second-best environments.

Beyond its theoretical contributions, the paper offers practical insights for economies with low rates of financial-sector participation and limited capacity to observe household income directly. The model suggests that governments can enhance welfare not necessarily by eliminating financial informality but by recognizing and designing around it. By embracing a more pragmatic and incentive-compatible tax policy, policymakers can sustain public good provision while minimizing distortions. Future work could extend this framework by incorporating political economy constraints or richer heterogeneity in preferences and enforcement mechanisms.

During the preparation of this work, the author used Claude and Refine.Ink for proof-check, grammar, spelling, and style corrections. The author reviewed and edited the output as needed and takes full responsibility for the content of the published article.

References

Albanesi, Stefania and Christopher Sleet, "Dynamic Optimal Taxation with Private Information," *Review of Economic Studies*, 2006, *73* (1), 1–30.

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman, "Tax Evasion and Inequality," *American Economic Review*, June 2019, *109* (6), 2073–2103.

Chodorow-Reich, Gabriel, Gita Gopinath, Prachi Mishra, and Abhinav Narayanan, "Cash and the Economy: Evidence from India's Demonetization," *The Quarterly Journal of Economics*, February 2020, *135* (1), 57–103.

Cole, Harold L. and Narayana R. Kocherlakota, "Efficient Allocations with Hidden Income and Hidden Storage," *Review of Economic Studies*, 2001, *68* (3), 523–542.

Farhi, Emmanuel and Xavier Gabaix, "Optimal Taxation with Behavioral Agents," *American Economic Review*, January 2020, *110* (1), 298–336.

Gelfand, I. M. and S. V. Fomin, *Calculus of Variations*, Englewood Cliffs, NJ: Prentice-Hall, 1963.

Golosov, Mikhail, Narayana R. Kocherlakota, and Aleh Tsyvinski, "Optimal Indirect and Capital Taxation," *Review of Economic Studies*, July 2003, *70* (3), 569–587.

Gordon, Roger and Wei Li, "Tax Structure in Developing Countries: Many Puzzles and a Possible Explanation," *Journal of Public Economics*, August 2009, *93* (7-8), 855–866.

Judd, Kenneth L., "The Law of Large Numbers with a Continuum of IID Random Variables," *Journal of Economic Theory*, 1985, *35* (1), 19–25.

Klapper, Leora, Dorothe Singer, Laura Starita, and Alexandra Norris, "The Global Findex Database 2025: Connectivity and Financial Inclusion in the Digital Economy," Technical Report, World Bank, Washington, DC 2025.

Kocherlakota, Narayana R., "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation," *Econometrica*, September 2005, *73* (5), 1587–1621.

_ , *The New Dynamic Public Finance* Toulouse Lectures in Economics, Princeton, NJ: Princeton University Press, 2010.

Mirrlees, James A., "An Exploration in the Theory of Optimum Income Taxation," *The Review of Economic Studies*, 1971, *38* (2), 175–208.

Myerson, Roger B., "Optimal Auction Design," *Mathematics of Operations Research*, 1981, *6* (1), 58–73.

Rogerson, William P., "The First-Order Approach to Principal-Agent Problems," *Econometrica*, November 1985, *53* (6), 1357–1367.

Samuelson, Paul A., "The Pure Theory of Public Expenditure," *The Review of Economics and Statistics*, 1954, *36* (4), 387–389.

Stiglitz, Joseph E., "The Causes and Consequences of the Dependence of Quality on Price," *Journal of Economic Literature*, 1987, *25* (1), 1–48.

Uhlig, Harald, "A Law of Large Numbers for Large Economies," *Economic Theory*, 1996, *8* (1), 41–50.

Vairo, Maren, "Robustly Optimal Income Taxation," *Working Paper*, 2025.

World Bank, "World Development Indicators," Technical Report, World Bank, Washington, DC 2024.

A Benchmark Model Proofs

A.1 Proof of Lemma 1

1. There is a unique $(b(e_k), 0)$ that solves agent e_k's problem.

Proof. First, since u is strictly concave and $h_1(b, b^c) = e_k - \tau(e) - b - b^c, h_2 = Rb + b^c + \psi g$ are linear functions of $(b, b^c$, then $f(b, b^c) = u(h_1(b, b^c)) + \beta u(h_2(b, b^c))$ is a strictly concave function of (b, b^c), as long as $R > 1$.[22] On top of this, since u is differentiable, then f is also differentiable, which implies that the first-order conditions are both necessary and sufficient.

Now, the Lagrangian of an agent with type e_k's problem (given the tax schedule) is:

$$\mathcal{L} = u\left(e_k - \tau(e_k) - b - b^c\right) + \beta u\left(Rb + b^c + \psi g\right) + \lambda_1 b + \lambda_2 b^c,$$

and this implies the following first-order conditions:

$$-u'\left(e_k - \tau(e_k) - b - b^c\right) + \beta R u'\left(Rb + b^c + \psi g\right) + \lambda_1 = 0, \tag{A1}$$

$$-u'\left(e_k - \tau(e_k) - b - b^c\right) + \beta u'\left(Rb + b^c + \psi g\right) + \lambda_2 = 0, \tag{A2}$$

$$b \geq 0, \;\; \lambda_1 \geq 0, \;\; \lambda_1 b = 0, \tag{A3}$$

$$b^c \geq 0, \;\; \lambda_2 \geq 0, \;\; \lambda_2 b^c = 0. \tag{A4}$$

Let's assume, for the sake of contradiction, that $b^c > 0$. Then, Equation (A4) implies that $\lambda_2 = 0$. If $b > 0$ then Equation (A3) implies that $\lambda_1 = 0$, which means that rewriting Equation (A1) and Equation (A2) we get:[23]

$$u'\left(e_k - \tau(e_k) - b(e_k) - b^c\right) = u'\left(Rb + b^c + \psi g\right),$$

$$u'\left(e_k - \tau(e_k) - b - b^c\right) = \beta u'\left(Rb + b^c + \psi g\right),$$

which then implies $\beta = 1$. This is not possible since $R > 1$ and $\beta = 1/R < 1$. So, the only other case is for $b = 0$. But then, Equation (A2) implies:

$$u'\left(e_k - \tau(e_k) - b^c\right) = \beta u'\left(b^c + \psi g\right),$$

which then follows $u'\left(e_k - \tau(e_k) - b^c\right) < u'\left(b^c + \psi g\right)$ since $\beta < 1$. But then, Equa-

22 Explicitly, $D^2 f = u''(h_1)\nabla h_1 \nabla h_1^\top + \beta u''(h_2)\nabla h_2 \nabla h_2^\top$, with $\nabla h_1 = (-1, -1)$ and $\nabla h_2 = (R, 1)$. Each term is a negative-semidefinite rank-one matrix, with null direction orthogonal to ∇h_1 and ∇h_2, respectively. Since ∇h_1 and ∇h_2 are linearly independent whenever $R \neq 1$ (guaranteed here by $R > 1$), no direction $z \neq 0$ is simultaneously orthogonal to both, so $z^\top D^2 f\, z < 0$ for every $z \neq 0$: $D^2 f$ is negative definite everywhere, and f is strictly concave in (b, b^c).

23 I am using here the assumption $\beta R = 1$.

tion (A1) implies that:

$$\lambda_1 = u'\left(e_k - \tau(e_k) - b^c\right) - u'\left(b^c + \psi g\right) < 0,$$

which contradicts Equation (A3). Hence, $b^c > 0$ cannot occur at the optimum, which means that $b^c = 0$. Also, due to this being a strict concave problem, the solution is unique, and it must be of the form $(b(e), 0)$.

□

2. If agent e_k has positive savings, then they are decreasing in $\tau(e_k)$ as well as in g.

Proof. If $b(e_k) > 0$, then $\lambda_1 = 0$ which implies that Equation (A1) can be rewritten as:[24]

$$u'\left(e - \tau(e) - b(e_k)\right) = u'\left(Rb(e_k) + \psi g\right).$$

Since u is strictly concave by Assumption 2, u' is strictly decreasing and therefore injective, so $e_k - \tau(e_k) - b = Rb + \psi g$. Solving for b we get:

$$b(e_k) = \frac{e_k - \tau(e_k) - \psi g}{1 + R}.$$

But since this could be negative, then $b(e) = max\{\frac{e_k - \tau(e_k) - \psi g}{1+R}, 0\}$. In the case of positive savings, we can see here that $b(e_k)$ is decreasing in $\tau(e_k)$ and g.

□

3. Every agent with positive savings consumes the same amount when young and old (although this may vary across agents).

Proof. If $b(e_k) > 0$, then Equation (A1) can be written as:[25]

$$u'\left(e_k - \tau(e_k) - b(e_k)\right) = u'\left(Rb(e_k) + \psi g\right).$$

And since u is strictly concave by Assumption 2, u' is strictly decreasing and therefore injective, which implies that consumption when old has to be equal to consumption when young.

□

A.2 Proof of Proposition 1

1. It is unique.

Proof. The problem I seek to solve is:

$$\max_{\{\tau(e_j)\}_j} \sum_{k=1}^{n} \left[u\left(e_k - \tau(e_k) - b(e_k)\right) + \beta u\left(Rb(e_k) + \psi g\right)\right] \pi_k,$$

24 Here I again use $\beta R = 1$ (Assumption 1): with $\lambda_1 = 0$, Equation (A1) reads $u'\left(e_k - \tau(e_k) - b(e_k)\right) = \beta R\, u'\left(Rb(e_k) + \psi g\right)$, which collapses to the expression below since $\beta R = 1$.

25 As in the proof of part 2, this step uses $\beta R = 1$ (Assumption 1) to simplify $\beta R\, u'\left(Rb(e_k) + \psi g\right)$ to $u'\left(Rb(e_k) + \psi g\right)$.

$$b(e_k) = \max\left\{\frac{e_k - \tau(e_k) - \psi g}{1+R}, 0\right\}, \quad g = \sum_{j=1}^{n} \tau(e_j)\pi_j.$$

Substituting these constraints into the objective, the problem reduces to maximizing $W(\{\tau(e_j)\}_j)$ over the unknowns $\{\tau(e_j)\}_j$. To establish strict concavity of W, note that for each k, the quantity

$$c_k \equiv \frac{e_k - \tau(e_k) - \psi \sum_j \tau(e_j)\pi_j}{1+R}$$

is an affine function of $\{\tau(e_j)\}_j$, so we consider two cases depending on the sign of c_k:

- When $c_k > 0$, so that $b(e_k) = c_k > 0$: substituting, young and old consumption both reduce to affine functions of $\{\tau(e_j)\}_j$. Since u is strictly concave by Assumption 2, each term $u(\cdot)$ in the sum is strictly concave in $\{\tau(e_j)\}_j$.
- When $c_k \leq 0$, so that $b(e_k) = 0$: young consumption is $e_k - \tau(e_k)$ and old consumption is $\psi \sum_j \tau(e_j)\pi_j$, both of which are affine in $\{\tau(e_j)\}_j$. Again, by Assumption 2, each term $u(\cdot)$ is strictly concave in $\{\tau(e_j)\}_j$.

In both cases each term in the sum is strictly concave, and the two expressions agree continuously at the boundary $c_k = 0$, so W is strictly concave globally. Since a positively weighted sum of strictly concave functions is strictly concave, and since the weights $\pi_k > 0$ for all k, W is strictly concave in $\{\tau(e_j)\}_j$. This guarantees a unique maximizer.

This argument uses $R < \psi$ strictly. If instead $\psi = R$, then for every k with $c_k > 0$, a common shift $\tau(e_j) \mapsto \tau(e_j) + s$ (applied to every type) raises g by s but leaves both c_k^y and c_k^o exactly unchanged, since the resulting change in $b(e_k)$ exactly offsets the change in $\tau(e_k)$ and ψg: this is a flat (non-strict) direction of W whenever the shift is small enough that no $b(e_k)$ crosses zero, and it is non-empty whenever $\mathcal{B} = \{1, \ldots, n\}$. Ruling out $\psi = R$ (Section 3) is exactly what excludes this degenerate direction and restores strict concavity of W along it. □

2. It equates consumption across all agents with positive savings, in the sense that agents with positive savings consume the same when young and old, and this consumption level is the same across all agents with positive savings.

   Proof. The problem I seek to solve is:

$$\max_{\{\tau(e_j)\}_j} \sum_{k=1}^{n} \left[u\left(e_k - \tau(e_k) - b(e_k)\right) + \beta u\left(Rb(e_k) + \psi g\right)\right]\pi_k,$$

$$b(e_k) = \max\left\{\frac{e_k - \tau(e_k) - \psi g}{1+R}, 0\right\}.$$

$$g = \sum_{j=1}^{n} \tau(e_j)\pi_j.$$

   Let $\mathcal{B} = \{k | b(e_k) > 0\}$. Then, using the expression for $b(e_k)$ together with the first-order conditions of agent k's problem, we get that if $k \in \mathcal{B}$:[26]

$$c_k^y = c_k^o = c_k = \frac{R(e_k - \tau(e_k)) + \psi \sum_{j=1}^{n} \tau(e_j)\pi_j}{1+R},$$

where c_k^y represents the optimal consumption of k when young, and c_k^o is the optimal consumption of agent k when old. For $k \notin \mathcal{B}$:

$$c_k^y = e_k - \tau(e_k), \;\; c_k^o = \psi \sum_{j=1}^{n} \tau(e_j)\pi_j.$$

Substituting this, then the objective function becomes:

$$\sum_{k \notin \mathcal{B}} \left[u\left(e_k - \tau(e_k)\right) + \beta u \left(\psi \sum_{j=1}^{n} \tau(e_j)\pi_j \right) \right] \pi_k +$$

$$\sum_{k \in \mathcal{B}} (1+\beta) u \left(\frac{R(e_k - \tau(e_k)) + \psi \sum_{j=1}^{n} \tau(e_j)\pi_j}{1+R} \right) \pi_k$$

The first-order condition (FOC) with respect to $\tau(e_j)$ for $j \notin \mathcal{B}$ is:

$$-u'(c_j^y)\pi_j + \sum_{k \notin \mathcal{B}} \beta\psi\pi_j u'(c_k^o)\pi_k + \sum_{k \in \mathcal{B}} \frac{\psi(1+\beta)\pi_j}{1+R} u'(c_k)\pi_k = 0,$$

while the FOC with respect to $\tau(e_j)$ if $j \in \mathcal{B}$ is:

$$-\frac{R(1+\beta)}{1+R} u'(c_j)\pi_j + \sum_{k \notin \mathcal{B}} \beta\psi\pi_j u'(c_k^o)\pi_k + \sum_{k \in \mathcal{B}} \frac{\psi(1+\beta)\pi_j}{1+R} u'(c_k)\pi_k = 0.$$

Manipulating these equations, we get:

$$u'(c_j^y) = \sum_{k \notin \mathcal{B}} \beta\psi u'(c_k^o)\pi_k + \sum_{k \in \mathcal{B}} \frac{\psi(1+\beta)}{1+R} u'(c_k)\pi_k, \;\; j \notin \mathcal{B},$$

$$u'(c_j) = \frac{1+R}{R(1+\beta)} \left[\sum_{k \notin \mathcal{B}} \beta\psi u'(c_k^o)\pi_k + \sum_{k \in \mathcal{B}} \frac{\psi(1+\beta)}{1+R} u'(c_k)\pi_k \right] \;\; j \in \mathcal{B}.$$

Notice that the right-hand side of the last expression is independent of j, which implies that if $j \in \mathcal{B}$ then $u'(c_j) = u'(c_k)$. And, since u is strictly increasing and concave, this implies that $c_j = c_k$ for all agents with positive savings.

□

3. If all agents have positive savings given $\tau^\star$, then the optimal tax schedule $\tau^\star$ is progressive.

26 Lemma 1 states that whenever agents have positive savings, they consume the same when young and old.

Proof. If all agents have positive savings then $\mathcal{B} = \{1, 2, ..., n\}$ and hence consumption (when young and old) is equal across all agents in the economy. But then, $c_j = c_k = c$ for $j \neq k$ implies, using the expression for c_k above, that $e_j - \tau(e_j) = e_k - \tau(e_k) = \frac{(1+R)c - \psi g}{R}$, i.e., disposable income $e_i - \tau(e_i)$ is the same constant across every agent with positive savings. Since $e_j - \tau(e_j) = e_k - \tau(e_k)$ for all j, k, it follows that $\tau(e_i) - \tau(e_{i-1}) = e_i - e_{i-1}$, so the marginal tax rate for all agents is one, which is a non-decreasing function. Hence, this tax schedule is progressive. □

4. The optimal tax schedule $\tau^\star$ implies zero financial informality and zero evasion.

   Proof. This is immediate by Lemma 1. □

B Hidden Savings Model Proofs

B.1 Proof of Lemma 2.

Throughout this proof, fix $e \in E$, $\tau \in \mathcal{T}$, and $g > 0$, and let $c_y = e - b - b^c - \tau(b)$ and $c_o = Rb + b^c + \psi g$ denote, respectively, consumption when young and old at the choice (b, b^c) under consideration.

1. The solution to $e \in E$'s problem is unique, and the first-order conditions are necessary and sufficient to characterize it.

   Proof. The problem that an agent $e \in E$ solves is:

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

   Let $F(b, b^c)$ denote the objective. Since F is continuous and the feasible set is compact,[27] a maximizer exists. I show it is unique and is exactly the point identified by the Karush–Kuhn–Tucker (KKT) conditions, *without* assuming F is jointly concave in (b, b^c): as discussed in the text following this lemma, that stronger property need not hold for every $\tau \in \mathcal{T}$, since nothing in Assumption 4 rules out $\tau''(b) < 0$.

   Two facts hold unconditionally, for every $\tau \in \mathcal{T}$ and every (b, b^c) in the feasible set:

$$F_{bb} = -u'(c_y)\tau''(b) + u''(c_y)(1 + \tau'(b))^2 + Ru''(c_o) < 0,$$

   using $\tau''(b) + (1 + \tau'(b))^2 \geq 0$ and $-u'(c) < u''(c)$ (Assumptions 3 and 4), and

$$F_{b^c b^c} = u''(c_y) + \beta u''(c_o) < 0,$$

   which holds for *any* τ at all, since it does not involve τ or its derivatives. That is, F is strictly concave along every line parallel to either axis, even at points where the full 2×2 Hessian may fail to be negative definite off those lines.

27 Feasibility requires consumption in both periods to remain in the domain of u, which together with $b, b^c \geq 0$ bounds the choice set.

Any maximizer must satisfy the KKT conditions derived in the proof of part 2 below, namely $-u'(c_y)(1+\tau'(b)) + u'(c_o) + \lambda_1 = 0$ and $-u'(c_y) + \beta u'(c_o) + \lambda_2 = 0$, with $\lambda_1, \lambda_2 \geq 0$, $\lambda_1 b = 0$, $\lambda_2 b^c = 0$. Complementary slackness places every KKT point in exactly one of four mutually exclusive cases:

(i) $b = b^c = 0$: a single candidate.

(ii) $b > 0, b^c = 0$: since $b \mapsto F(b,0)$ is strictly concave ($F_{bb} < 0$ above), it has at most one maximizer on $[0, e_n]$, so at most one candidate.

(iii) $b = 0, b^c > 0$: symmetrically, $b^c \mapsto F(0, b^c)$ is strictly concave ($F_{b^c b^c} < 0$), so at most one candidate.

(iv) $b > 0, b^c > 0$: then $\lambda_1 = \lambda_2 = 0$, and both first-order conditions hold with equality, which (as shown in part 4 below) forces $\tau'(b) = R - 1$.

For case (iv), let $\mathcal{B} := \{b \in (0, e_n] : \tau'(b) = R - 1\}$. If $\mathcal{B} = \emptyset$, this case contributes nothing. Otherwise fix any $b_0 \in \mathcal{B}$: along the line $b = b_0$, c_y and c_o move affinely and in opposite directions as b^c varies, so the condition $u'(c_o) = R\,u'(c_y)$ equivalent to both first-order conditions holding jointly at b_0 has at most one solution b^c, since its derivative in b^c is $u''(c_o) + Ru''(c_y) < 0$. This gives a unique candidate $(b_0, b^c(b_0))$ for each $b_0 \in \mathcal{B}$, and since $F_b = F_{b^c} = 0$ at every such point, F takes the *same* value across all of them: moving along $\mathcal{B}$, $dF = F_b\,db + F_{b^c}\,db^c = 0$ throughout. So case (iv) contributes at most one candidate value.

Finally, a case (iv) candidate is only a genuine local (hence possible global) maximum if the Hessian is negative semidefinite there. Direct computation of the Hessian determinant at $(b_0, b^c(b_0))$ gives[28] $u'(c_y)\,\tau''(b_0) \cdot Q$ for $Q := -u''(c_y) - u''(c_o)/R > 0$, which is ≤ 0 if and only if $\tau''(b_0) \leq 0$. Hence a case (iv) candidate is a saddle point – not eligible for the comparison below – whenever $\tau''(b_0) < 0$, and is retained only when $\tau''(b_0) = 0$.

Collecting cases (i)–(iv), the agent's problem has a finite list of at most four candidate values, and its global maximum (which exists, by compactness) equals the largest of these. This establishes that the first-order conditions are necessary and sufficient to characterize the solution – every maximizer is one of these KKT-consistent candidates, and every such candidate has been checked – and uniqueness: whenever two candidates tie, they attain the same maximized utility, and this coincidence does not affect any downstream argument, since those depend only on the maximized utility level or on which case is realized (itself uniquely determined even when, within case (iv), the split between b and b^c is not – see the discussion following Lemma B.1). □

2. If $b_{\tau,g}(e) > 0$ and $b^c_{\tau,g}(e) = 0$ then:

$$\tau'\left(b_{\tau,g}(e)\right) \leq R - 1.$$

Proof. The Lagrangian of agent e's problem is:

$$\mathcal{L} = u\left(e - b - b^c - \tau(b)\right) + \beta u\left(Rb + b^c + \psi g\right) + \lambda_1 b + \lambda_2 b^c,$$

28 Writing $a := u''(c_y)$, $o := u''(c_o)$, $p := u'(c_y)$, $m := 1 + \tau'(b)$: $F_{bb}F_{b^c b^c} - F^2_{bb^c} = ao(m-R)^2/R + p\,\tau''(b)\,Q$ for $Q := -a - o/R > 0$; on $\mathcal{B}$, $m = R$, so the first term vanishes.

which yields the first-order (Karush–Kuhn–Tucker) conditions that fully characterize agent e's problem:

$$-u'(c_y)(1+\tau'(b)) + \beta R\, u'(c_o) + \lambda_1 = 0,$$
$$-u'(c_y) + \beta u'(c_o) + \lambda_2 = 0,$$
$$\lambda_1 \geq 0,\;\; b \geq 0,\;\; \lambda_1 b = 0,$$
$$\lambda_2 \geq 0,\;\; b^c \geq 0,\;\; \lambda_2 b^c = 0.$$

Using $\beta R = 1$ (Assumption 1), the first condition can be rewritten as $u'(c_y)(1+\tau'(b)) = u'(c_o) + \lambda_1$, and the second as $u'(c_y) = \beta u'(c_o) + \lambda_2$. These are the conditions used for the remainder of this proof (parts 2–4).

If $b > 0$ then $\lambda_1 = 0$. The FOC with respect to b can be rewritten as:

$$1 + \tau'(b) = \frac{u'(c_o)}{u'(c_y)},$$

which then, using the FOC of b^c and $\lambda_2 \geq 0$:

$$1 + \tau'(b) = \frac{u'(c_o)}{u'(c_y)} \leq \frac{1}{\beta} = R,$$

and hence $\tau'(b) \leq R - 1$.

□

3. If $b_{\tau,g}(e) = 0$ and $b^c_{\tau,g}(e) > 0$ then:

$$\tau'(0) \geq R - 1.$$

Proof. Again, using the FOCs derived in the proof of part 2, since $b^c > 0$ we conclude that $\lambda_2 = 0$ and that:

$$R = \frac{1}{\beta} = \frac{u'(c_o)}{u'(c_y)}.$$

Combining this with the FOC with respect to b we have:

$$R = \frac{u'(c_o)}{u'(c_y)} \leq 1 + \tau'(0),$$

which yields $\tau'(0) \geq R - 1$. □

4. If $b_{\tau,g}(e) > 0$ and $b^c_{\tau,g}(e) > 0$ then:

$$\tau'\left(b_{\tau,g}(e)\right) = R - 1.$$

Proof. Finally, if both $b, b^c > 0$ then $\lambda_1 = \lambda_2 = 0$, and then we have:

$$u'(c_y)(1+\tau'(b)) = u'(c_o), u'(c_y) = \beta u'(c_o),$$

which implies that $1 + \tau'(b) = 1/\beta = R$.

□

B.2 Properties of $\mathcal{G}(\tau)$.

Throughout this section, I consider the following norm that induces a distance in $\mathcal{T}$.

Definition 4. Consider $\mathbb{B} = [0, e_n]$. Let $||\cdot|| : \mathcal{T} \to \mathbb{R}$ be given by:

$$||\tau|| = \max_{b \in \mathbb{B}} |\tau(b)| + \max_{b \in \mathbb{B}} |\tau'(b)| + \max_{b \in \mathbb{B}} |\tau''(b)|$$

This, as shown in Gelfand and Fomin (1963) is a norm in the space of twice continuously differentiable functions. The induced distance between two different functions is $||\tau - \tilde{\tau}||$. Notice that if $||\tau - \tilde{\tau}|| < \epsilon$ then:

$$|\tau(b) - \tilde{\tau}(b)| < \epsilon \quad |\tau'(b) - \tilde{\tau}'(b)| < \epsilon \quad |\tau''(b) - \tilde{\tau}''(b)| < \epsilon,$$

for all $b \in \mathbb{B}$, and hence, two functions are "close" according to this norm if the function and its first and second derivatives are "close".

Lemma B.1. *The correspondence $\mathcal{G} : \mathcal{T} \rightrightarrows \mathbb{R}_+$ is non-empty and it is actually a function.*

Proof. Let $\tau \in \mathcal{T}$, $g > 0$, $e \in E$. First, I show that $(b_{\tau,g}(e), b^c_{\tau,g}(e))$ are a function of g and that they are both continuous in g. By Lemma 2 (parts 2–4), the solution satisfies $\min\{b_{\tau,g}(e), b^c_{\tau,g}(e)\} = 0$ except at the knife-edge case $\tau'(b_{\tau,g}(e)) = R - 1$ with $b_{\tau,g}(e) > 0$, under which both can be positive but only the sum $Rb_{\tau,g}(e) + b^c_{\tau,g}(e)$ is pinned down (this non-generic case is excluded from the argument below). Away from this knife edge, the agent is either bank-only ($b > 0, b^c = 0$) or couch-only ($b = 0, b^c > 0$), and I analyze the comparative statics of each regime in turn, before checking that the two pieces agree at the boundary between them.

Bank-only regime ($b > 0$, $b^c = 0$). Here b^c is fixed at its corner value 0, and the agent solves the one-variable problem $\max_b \, u\left(e - b - \tau(b)\right) + \beta u\left(Rb + \psi g\right)$, whose first-order condition is:[29]

$$F_\tau(b, g) = -u'(c_y)(1 + \tau'(b)) + u'(c_o) = 0.$$

Hence:

$$F_b(b, g) = -u'(c_y)\tau''(b) + u''(c_y)(1 + \tau'(b))^2 + Ru''(c_o),$$

$$F_g(b, g) = \psi u''(c_o).$$

Since $\tau \in \mathcal{T}$ and $-u'(c) \le u''(c) < 0$ then:

$$F_b(b, g) \le u''(c_y)\left[\tau''(b) + (1 + \tau'(b))^2\right] + Ru''(c_o) < 0, \tag{B5}$$

$$F_g(b, g) = \psi u''(c_o) < 0. \tag{B6}$$

Hence, since $F_b \neq 0$ then by the implicit function theorem:

$$\frac{\partial b}{\partial g} = -\frac{F_g}{F_b} < 0,$$

29 As before, this expression already uses $\beta R = 1$ (Assumption 1) to replace $\beta R\, u'(c_o)$ with $u'(c_o)$.

which implies two things: first that b is differentiable with respect to g (and hence b must be continuous with respect to g); and that b is strictly decreasing in g. Let $b^u_{\tau,g}(e)$ be the (unconstrained, within this regime) solution. The bank-only branch of the solution to the constrained problem is then:

$$b_{\tau,g}(e) = \max\left\{b^u_{\tau,g}(e), 0\right\},$$

which is also continuous, and differentiable almost for all g except at the $\tilde{g}$ such that $b^u_{\tau,\tilde{g}}(e) = 0$ (which is unique since this is a strictly decreasing function of g). Figure X plots this construction.

Couch-only regime ($b = 0,\ b^c > 0$). Symmetrically, here b is fixed at its corner value 0, and the agent solves the one-variable problem $\max_{b^c}\ u\left(e - \tau(0) - b^c\right) + \beta u\left(b^c + \psi g\right)$, whose first-order condition is:

$$\tilde{F}(b^c, g) = -u'(c_y) + \beta u'(c_o) = 0, \qquad c_y = e - \tau(0) - b^c, \quad c_o = b^c + \psi g.$$

Hence $\tilde{F}_{b^c}(b^c, g) = -u''(c_y) - \beta u''(c_o) < 0$, since both terms are strictly positive under Assumption 2 ($u'' < 0$), and $\tilde{F}_g(b^c, g) = \beta\psi u''(c_o) < 0$. Since $\tilde{F}_{b^c} \neq 0$, the implicit function theorem gives:

$$\frac{\partial b^c}{\partial g} = -\frac{\tilde{F}_g}{\tilde{F}_{b^c}} < 0,$$

so b^c is differentiable and strictly decreasing in g within this regime, by exactly the same argument as the bank-only case. Writing $b^{c,u}_{\tau,g}(e)$ for the (unconstrained, within this regime) solution, the couch-only branch is $b^c_{\tau,g}(e) = \max\{b^{c,u}_{\tau,g}(e), 0\}$, which is continuous by the same reasoning as above.

Agreement at the regime boundary. By Lemma 2, an interior bank-only optimum requires $\tau'(b) \leq R - 1$ at the chosen $b > 0$, while an interior couch-only optimum requires $\tau'(0) \geq R - 1$; the two regimes can only be adjacent to each other through the shared corner $b = b^c = 0$, at which both branches above agree by construction (each equals $\max\{\cdot, 0\}$ of a continuous unconstrained candidate). Hence a transition between regimes as g varies passes continuously through $b = b^c = 0$, rather than jumping directly from positive bank savings to positive couch savings or vice versa, and the pieced-together pair $(b_{\tau,g}(e), b^c_{\tau,g}(e))$ is a continuous, non-increasing function of g throughout.[30]

Let $G_\tau(g)$ be defined as follows:

$$G_\tau(g) = \sum_{i=1}^{n} \tau\left(b_{\tau,g}(e_i)\right)\pi_i - g = PGC_\tau(g) - g.$$

Since τ is continuous in b and $b_{\tau,g}(e)$ is continuous in g for all e_i, then this function is continuous in g. Furthermore, $G(0) \geq 0$ since the government is constrained to provide a non-negative amount of the public good, and $PCG(g)$ is bounded, since τ is a bounded function. Hence $\lim_{g\to\infty} G_\tau(g) = -\infty$. Hence, there exists $\tilde{g} > 0$ such that $G(\tilde{g}) < 0$. In conclusion, $G(g)$ is continuous, $G(0) \geq 0$, and $G(\tilde{g}) < 0$. Then, by Bolzano's Theorem, there must exist at least

30 This construction pins down $(b_{\tau,g}(e), b^c_{\tau,g}(e))$ within whichever regime Lemma 2 identifies as containing the agent's necessary (KKT) optimum. The proof of Lemma 2 (part 1) rules out a higher-utility optimum in any other regime directly, by comparing across the (at most four) KKT-consistent candidates – it does not require, and does not assume, that the agent's problem is jointly concave in (b, b^c).

Figure X: CONTINUITY OF $b_{\tau,g}(e)$ WITH RESPECT TO g.

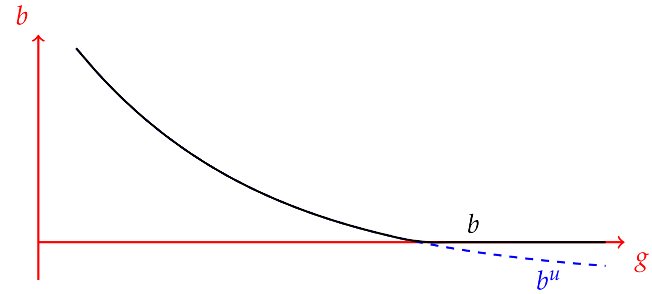

one $\hat{g} \in [0, \tilde{g})$ such that $G(\hat{g}) = 0$. Notice that $\hat{g} \in \mathcal{G}(\tau)$, and hence this set is non-empty.

I now show that this is a function of τ. For this, I need to show that $G_\tau(g)$ can only cross the y-axis once. To do this, I show that this function is strictly decreasing in g. The chain rule implies:

$$\frac{\partial G_\tau}{\partial g} = -1 + \sum_{i=1}^{n} \pi_i \tau' \left(b_{\tau,g}(e_i) \right) \frac{b_{\tau,g}(e_i)}{\partial g}.$$

I now consider three cases:

1. **Case 1:** τ is strictly increasing.

   In this case, since $\partial b / \partial g < 0$ and $\tau'(b) > 0$, then for all g:

   $$\frac{\partial G_\tau}{\partial g} < 0,$$

   which implies that this function is strictly decreasing in g.

2. **Case 2:** τ is strictly decreasing.

   Using Equation (B5) and Equation (B6) together with Assumption 3 we get:

   $$\left| \frac{b_{\tau,g}(e_i)}{\partial g} \right| \leq \frac{|\psi u''(c_o)|}{|u'(c_y)[\tau''(b) + (1+\tau'(b))^2] - R\psi u''(c_o)|}.$$

   Since Assumption 4 implies $\tau''(b) + (1+\tau'(b))^2 \geq 0$, $u' > 0$ and $u'' < 0$ then:

   $$\left| \frac{b_{\tau,g}(e_i)}{\partial g} \right| \leq \frac{|\psi u''(c_o)|}{| - R\psi u''(c_o)|} = \frac{1}{R} < 1.$$

   Since $0 < |\tau'(b)| < 1$ by Assumption 4 (and because τ is decreasing) then each term $\tau' \left(b_{\tau,g}(e_i) \right) \frac{b_{\tau,g}(e_i)}{\partial g}$ can be at most one (but not exactly equal to one), which means that:

   $$\frac{\partial G_\tau}{\partial g} = -1 + \sum_{i=1}^{n} \pi_i \tau' \left(b_{\tau,g}(e_i) \right) \frac{b_{\tau,g}(e_i)}{\partial g} < 0.$$

3. **Case 3:** τ' changes sign across $b \in [0, e_n]$.

   Partition $\{1, \ldots, n\}$ into $\mathcal{I}^+ = \{i : \tau'(b_{\tau,g}(e_i)) \geq 0\}$ and $\mathcal{I}^- = \{i : \tau'(b_{\tau,g}(e_i)) < 0\}$. For $i \in \mathcal{I}^+$, since $\partial b_{\tau,g}(e_i)/\partial g < 0$ (established above, independently of the sign of τ') and $\tau'(b_{\tau,g}(e_i)) \geq 0$, the term $\tau'(b_{\tau,g}(e_i))\, \partial b_{\tau,g}(e_i)/\partial g$ is non-positive. For $i \in \mathcal{I}^-$, both factors are negative, so the term is positive; moreover, the bound $\left|\partial b_{\tau,g}(e_i)/\partial g\right| \leq 1/R$ established in Case 2 above did not rely on τ being globally decreasing (it follows directly from Equations (B5) and (B6), which hold for every $\tau \in \mathcal{T}$), so it applies here as well. Combined with $|\tau'(b_{\tau,g}(e_i))| < 1$, this gives, for $i \in \mathcal{I}^-$:

$$0 < \tau'(b_{\tau,g}(e_i)) \frac{\partial b_{\tau,g}(e_i)}{\partial g} = \left|\tau'(b_{\tau,g}(e_i))\right| \left| \frac{\partial b_{\tau,g}(e_i)}{\partial g} \right| < 1 \cdot \frac{1}{R} = \frac{1}{R}.$$

   Therefore:

$$\sum_{i=1}^{n} \pi_i \tau'(b_{\tau,g}(e_i)) \frac{\partial b_{\tau,g}(e_i)}{\partial g} = \sum_{i \in \mathcal{I}^+} \pi_i \underbrace{\tau'(b_{\tau,g}(e_i)) \frac{\partial b_{\tau,g}(e_i)}{\partial g}}_{\leq 0} + \sum_{i \in \mathcal{I}^-} \pi_i \underbrace{\tau'(b_{\tau,g}(e_i)) \frac{\partial b_{\tau,g}(e_i)}{\partial g}}_{< 1/R}$$

$$< \frac{1}{R} \sum_{i \in \mathcal{I}^-} \pi_i \leq \frac{1}{R} \sum_{i=1}^{n} \pi_i = \frac{1}{R} < 1,$$

   so that $\dfrac{\partial G_\tau}{\partial g} = -1 + \sum_{i=1}^{n} \pi_i \tau'(b_{\tau,g}(e_i)) \dfrac{\partial b_{\tau,g}(e_i)}{\partial g} < -1 + \dfrac{1}{R} < 0$. Hence $G_\tau(g)$ is strictly decreasing in g in this case as well, which, together with Cases 1 and 2, establishes strict monotonicity of $G_\tau(g)$ for every $\tau \in \mathcal{T}$.

$\square$

Lemma B.2. *Consider the $\mathcal{C}_2$-norm over $\mathcal{T}$, $||\cdot||$. The following continuity properties hold:*

1. *$\mathcal{G}(\cdot)$ is a $||\cdot||$-continuous function.*

   Proof. I first, show that for every $g > 0$ the functions $G_\tau(g) = PGC_\tau(g) - g$ and $G_{\tilde{\tau}}(g)$ (defined analogously) are "close" whenever $\tau, \tilde{\tau}$ are "close" according to the $||\cdot||$ distance. Given a fixed $g > 0$, the FOC that implicitly characterizes $b_{\tau,g}$ is a function of τ and τ'. Then, if $\tau, \hat{\tau}$ are close under the $||\cdot||$ norm, then both the function and their respective derivatives are (uniformly) close, implying that $b_{\tau,g}$ is $||\cdot||-$continuous function.

   Since $PGC_\tau(g) = \sum_{i=1}^{n} \tau\left(b_{\tau,g}(e_i)\right) \pi_i$ and τ itself is continuous in its argument, then PGC_τ and therefore $G_\tau(g)$ are continuous in τ.

   Now, let's assume that the function $\mathcal{G}(\tau)$ is not continuous, meaning that the equilibrium value of public good provision changes discontinuously with the tax schedule. This means that for some τ there is $\epsilon > 0$ such that for every $\delta > 0$ there is $\tilde{\tau}$ such that $||\tau - \tilde{\tau}|| < \delta$ but $|\mathcal{G}(\tau) - \mathcal{G}(\tilde{\tau})| \geq \epsilon$. Without loss, let's assume that $\mathcal{G}(\tau) < \mathcal{G}(\tilde{\tau})$ (i.e., there is a larger public good provision in equilibrium under $\tilde{\tau}$). Then, as Figure XI shows, there must be (at least) an ϵ distance between $\mathcal{G}(\tau)$ and $\mathcal{G}(\tilde{\tau})$. Also, since $G_\tau(\cdot)$

is strictly decreasing, then $G_\tau(\mathcal{G}(\tilde{\tau})) < 0$, which means that there is $\eta > 0$ such that $(G_\tau(\mathcal{G}(\tilde{\tau}) + \eta, G_\tau(\mathcal{G}(\tilde{\tau}) - \eta) \subset \mathbb{R} \setminus [0, \infty)$. But then, since $G_\tau(g)$ is a continuous function, given $\eta > 0$ there must exist $\delta > 0$ such that $G_{\tilde{\tau}}(\mathcal{G}(\tilde{\tau}))$ in within this open interval, which contains all negative numbers. But this cannot be since $\mathcal{G}(\tilde{\tau})$ is a zero of $G_{\tilde{\tau}}(\cdot)$. Figure XI presents a graph showing this argument. In conclusion, $\mathcal{G}(\tau)$ is continuous.

Figure XI: $\mathcal{G}(\tau)$ IS CONTINUOUS.

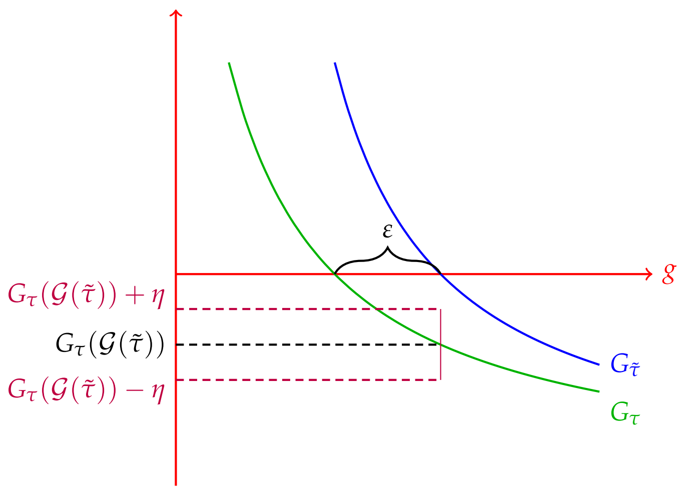

$\square$

2. *The function $b_\cdot(\tilde{e}) : \mathcal{T} \to \mathbb{R}_+$ given by:*

$$b_\cdot(\tilde{e}) = b_{\cdot,\mathcal{G}(\cdot)}(\tilde{e}),$$

is a $||\cdot||$-continuous for all $\tilde{e} \in E$.

Proof. If we replace g for $\mathcal{G}(\tau)$ in the first-order conditions that characterize the optimal savings, these become a continuous function of τ and τ'. But then, any other function $\hat{\tau}$ that is close in the $\mathcal{C}_2$-way will imply an optimal savings $b_{\hat{\tau},\mathcal{G}(\hat{\tau})}$ that is close to $b_{\tau,\mathcal{G}(\tau)}$, implying this is continuous in τ. $\square$

B.3 Proof of Lemma 3.

Proof. I prove only the case $b > 0, b^c = 0$, the rest can be shown using exactly the same arguments with some minor adjustments. The relaxed-program that agent e solves is:

$$\max_b \; u\left(e - b - \tau(b)\right) + \beta u\left(Rb + \psi g\right).$$

The first order condition that characterizes the optimal solution to this problem is:[31]

$$F(b, e, g) = -u'(c_y)(1 + \tau'(b)) + u'(c_o) = 0.$$

31 This uses $\beta R = 1$ (Assumption 1) to replace $\beta R\, u'(c_o)$ with $u'(c_o)$.

Hence:

$$F_b = -u'(c_y)\tau''(b) + u''(c_y)(1+\tau'(b))^2 + Ru''(c_o),$$
$$F_e = -u''(c_y)(1+\tau'(b)),$$
$$F_g = \psi u''(c_o).$$

Given Assumption 3 and Assumption 4, and arguing as in the proof of Lemma 2 (part 1), $F_b < u''(c_y)\left[\tau''(b) + (1+\tau'(b))^2\right] \leq 0$. Also, since $u'' < 0$ by Assumption 2, $F_g = \psi u''(c_o) < 0$ always. This implies:

$$\frac{\partial b^u}{\partial g} = -\frac{F_g}{F_b} < 0,$$

since $-F_g > 0$ and $F_b < 0$, so their ratio is negative.
Now, since $u'' < 0$, $F_e = -u''(c_y)(1+\tau'(b))$ has the same sign as $1+\tau'(b)$, so $F_e > 0$ if and only if $1+\tau'(b) > 0$; note that this holds for every $b \in [0, e_n]$ except possibly at points where $\tau'(b) = -1$, since Assumption 4 requires $\tau'(b) \geq -1$. Whenever $F_e > 0$:

$$\frac{\partial b^u}{\partial e} = -\frac{F_e}{F_b} > 0,$$

since $-F_e < 0$ and $F_b < 0$, so their ratio is positive.
To go from the unconstrained case to the actual agent's problem, I follow the same construction that I did in Section B.2, which guarantees that b is a continuous non-increasing function of g and a continuous non-decreasing function of e.

□

B.4 Proof of Lemma 4.

1. If $b_{\tau,g}(e) > 0$ then:

   (a) $c^y_{\tau,g}(e) > c^o_{\tau,g}(e)$ if and only if $\tau'(b_{\tau,g}(e)) > 0$.

   Proof. The first order condition when $b > 0$ is:[32]

   $$u'(c_y)(1+\tau'(b)) = u'(c_o).$$

   Suppose $\tau'(b) > 0$. Since $u'(c_y) > 0$ (Assumption 2) and $1+\tau'(b) > 1$, we get $u'(c_o) = u'(c_y)(1+\tau'(b)) > u'(c_y)$, i.e. $u'(c_y) < u'(c_o)$. Since u is strictly concave, u' is strictly decreasing, so $u'(c_y) < u'(c_o)$ implies $c_y > c_o$.
   Conversely, suppose $c_y > c_o$. Since u' is strictly decreasing and strictly positive, $0 < u'(c_y) < u'(c_o)$, so the first-order condition gives:

   $$1+\tau'(b) = \frac{u'(c_o)}{u'(c_y)} > 1,$$

   yielding $\tau'(b) > 0$. □

   (b) $c^y_{\tau,g}(e) < c^o_{\tau,g}(e)$ if and only if $\tau'(b_{\tau,g}(e)) < 0$.

32 As before, this uses $\beta R = 1$ (Assumption 1) to replace $\beta R\, u'(c_o)$ with $u'(c_o)$.

Proof. Suppose $\tau'(b) < 0$. By Assumption 4, $\tau'(b) \geq -1$, so $0 \leq 1 + \tau'(b) < 1$. Since $u'(c_y) > 0$, the first-order condition gives $u'(c_o) = u'(c_y)(1 + \tau'(b)) < u'(c_y)$, i.e. $u'(c_o) < u'(c_y)$, which since u' is strictly decreasing implies $c_o > c_y$.

Conversely, suppose $c_y < c_o$. Since u' is strictly decreasing and strictly positive, $0 < u'(c_o) < u'(c_y)$, so the first-order condition gives $1 + \tau'(b) = u'(c_o)/u'(c_y) \in (0, 1)$, yielding $\tau'(b) < 0$. □

(c) $c^y_{\tau,g}(e) = c^o_{\tau,g}(e)$ if and only if $\tau'(b_{\tau,g}(e)) = 0$.

Proof. If $\tau'(b) = 0$, the first-order condition gives $u'(c_o) = u'(c_y)$, and since u' is injective (being strictly decreasing), $c_o = c_y$. Conversely, if $c_y = c_o$, then $u'(c_y) = u'(c_o)$, so the first-order condition gives $1 + \tau'(b) = 1$, i.e. $\tau'(b) = 0$. □

2. If $b^c_{\tau,g}(e) > 0$ then $c^y_{\tau,g}(e) > c^o_{\tau,g}(e)$.

Proof. If $b^c > 0$ then the first-order condition with respect to b^c is:

$$u'(c_y) = \beta u'(c_o).$$

Since $R > 1$ then $\beta = 1/R < 1$, implying $u'(c_y) < u'(c_o)$, which by Assumption 2 leads to $c_y > c_o$.

□

B.5 Proof of Proposition 2.

Proof. Suppose $\tau \in \mathcal{T}$ is such that $b_{\tau,g}(e) > 0$ and $b^c_{\tau,g}(e) > 0$ for some $e \in E$, with $g = \mathcal{G}(\tau)$. I construct $\hat{\tau} \in \mathcal{T}$ that induces a strict Pareto improvement over τ, which implies τ cannot be optimal.

Construction. Write $b^\star := b_{\tau,g}(e)$. Since $b^\star > 0$ and $b^c_{\tau,g}(e) > 0$, Lemma 2 gives $\tau'(b^\star) = R - 1$. Because E is finite, fix $w > 0$ smaller than the distance from $b^\star$ to every other agent's equilibrium bank-saving level $b_{\tau,g}(e_i)$, $e_i \neq e$ (if some other agent shares the same level $b^\star$, group her together with e in what follows, since she reacts in the same direction). For $\delta \in (0, R - 1)$, let $\hat{\tau} \in \mathcal{T}$ agree with τ outside $(b^\star, b^\star + w)$ – so that $\hat{\tau}$ and τ, and their first and second derivatives, coincide at $b^\star$ and to its left – and be strictly flatter than τ on $(b^\star, b^\star + w)$, with $\hat{\tau}'(b) = R - 1 - \delta$ on the bulk of this interval, smoothly tapering back to $\tau'(b^\star + w)$ at the right endpoint via a standard $\mathcal{C}^2$ interpolation (e.g., a cubic polynomial patch matching the prescribed value and slope at both ends of the tapering region); for δ, w small enough this remains in $\mathcal{T}$. In words, $\hat{\tau}$ is pinned to τ's value at $b^\star$ but strictly lowers the marginal tax rate just above it, leaving the schedule exactly unchanged everywhere else, including near every other agent's equilibrium choice.

Step 1: agent e strictly increases her bank savings, at the original g. I show directly that $\hat{b} := b_{\hat{\tau},g}(e) > b^\star$, without appealing to joint concavity of her problem (which, as discussed after Lemma 2, need not hold at a point with both $b, b^c > 0$).

First, since $\hat{\tau} = \tau$ on $[0, b^\star]$, the agent's objective $F_{\hat{\tau}}$ and F_τ coincide at every (b, b^c) with $b \leq b^\star$; since $(b^\star, b^c_{\tau,g}(e))$ maximizes F_τ over the whole feasible set (in particular over every point with $b \leq b^\star$), it follows that $(b^\star, b^c_{\tau,g}(e))$ weakly dominates, under $\hat{\tau}$ as well, every feasible point with $b \leq b^\star$.

Second, I exhibit a specific feasible point with $b > b^\star$ that strictly beats $(b^\star, b^c_{\tau,g}(e))$ under $\hat{\tau}$. Write $b^{c\star} := b^c_{\tau,g}(e) > 0$ and consider the line $\ell(t) := (b^\star + t,\, b^{c\star} - t)$ for $t \in [0, b^{c\star})$, which shifts a dollar from couch into bank savings while holding $b + b^c$ fixed. Along this line, $c_y(t) = e - (b^\star + b^{c\star}) - \hat{\tau}(b^\star + t)$ and $c_o(t) = Rb^\star + b^{c\star} + \psi g + (R-1)t$, so

$$\frac{d}{dt} F_{\hat{\tau}}(\ell(t)) = -u'(c_y(t))\, \hat{\tau}'(b^\star + t) + \beta(R-1)\, u'(c_o(t)).$$

At $t = 0$: $\hat{\tau}'(b^\star) = \tau'(b^\star) = R - 1$ (the construction matches derivatives at $b^\star$), and the original first-order condition for b^c at $(b^\star, b^{c\star})$ under τ gives $\beta u'(c_o(0)) = u'(c_y(0))$; substituting, $\frac{d}{dt} F_{\hat{\tau}}(\ell(t))\big|_{t=0} = (R-1)\left[\beta u'(c_o(0)) - u'(c_y(0))\right] = 0$, as expected at a critical point. For $t > 0$ small, $\hat{\tau}'(b^\star + t) = R - 1 - \delta$ (the flattened region), so

$$\frac{d}{dt} F_{\hat{\tau}}(\ell(t)) = \delta\, u'(c_y(t)) + (R-1)\left[\beta u'(c_o(t)) - u'(c_y(t))\right].$$

The bracketed term is continuous and vanishes at $t = 0$, while $u'(c_y(t)) \to u'(c_y(0)) > 0$; hence the right-hand side converges to $\delta\, u'(c_y(0)) > 0$ as $t \to 0^+$. By continuity, there exists $t_0 \in (0, b^{c\star})$ such that $\frac{d}{dt} F_{\hat{\tau}}(\ell(t)) > 0$ on $(0, t_0]$, so $F_{\hat{\tau}}(\ell(t_0)) > F_{\hat{\tau}}(\ell(0)) = F_{\hat{\tau}}(b^\star, b^{c\star})$.

Combining the last two paragraphs: the point $\ell(t_0)$, with b-component $b^\star + t_0 > b^\star$, strictly beats $(b^\star, b^{c\star})$ under $\hat{\tau}$, while every point with $b \leq b^\star$ is at best tied with $(b^\star, b^{c\star})$ under $\hat{\tau}$. So the agent's maximizer under $(\hat{\tau}, g)$ cannot have $b \leq b^\star$, i.e. $\hat{b} > b^\star$.

Step 2: no other agent's choice is affected, and revenue collected from e strictly rises, at the original g. By Lemma 2, agent e_i's problem under τ has a unique maximizer at $b_i^\star := b_{\tau,g}(e_i)$; since F_τ is continuous and $b_i^\star$ is the unique maximizer, F_τ is bounded strictly below $F_\tau(b_i^\star, b_i^{c\star})$ everywhere on the compact set of feasible points with b-component inside $(b^\star, b^\star + w)$, by an amount that (as w shrinks) depends only on the (fixed, positive) distance from $b_i^\star$ to this interval, while the largest tax saving available anywhere inside the interval is at most of order $\delta \cdot w$. Hence, for δ, w small enough, no other agent finds it optimal to deviate into $(b^\star, b^\star + w)$, and her choice under $(\hat{\tau}, g)$ coincides with $(b_i^\star, b_i^{c\star})$. Since $\hat{\tau}(b^\star) = \tau(b^\star)$ and $\hat{\tau}'(b) = R - 1 - \delta > 0$ for b near $b^\star$, $\hat{\tau}$ is strictly increasing immediately to the right of $b^\star$, so $\hat{\tau}(\hat{b}) > \hat{\tau}(b^\star) = \tau(b^\star)$. Therefore

$$\mathrm{PGC}_{\hat{\tau}}(g) = \pi_e \hat{\tau}(\hat{b}) + \sum_{e_i \neq e} \pi_i \tau(b_i^\star) \;>\; \pi_e \tau(b^\star) + \sum_{e_i \neq e} \pi_i \tau(b_i^\star) \;=\; \mathrm{PGC}_\tau(g) \;=\; g,$$

using budget balance at the original equilibrium in the last equality. Intuitively, e was exactly indifferent, at the margin, between the bank and the couch; inducing her to shift a marginal dollar from the untaxed couch to the taxed bank is a pure revenue gain that she does not internalize in her own decision.

Step 3: the public good strictly rises, $\mathcal{G}(\hat{\tau}) > g$. By Lemma B.1, $G_{\hat{\tau}}(g') := \mathrm{PGC}_{\hat{\tau}}(g') - g'$ is continuous and strictly decreasing in g', with a unique root $\mathcal{G}(\hat{\tau})$, since $\hat{\tau} \in \mathcal{T}$. Step 2 shows $G_{\hat{\tau}}(g) > 0$; since $G_{\hat{\tau}}$ is strictly decreasing, its unique root must exceed g, i.e., $\hat{g} := \mathcal{G}(\hat{\tau}) > g$.

Step 4: strict Pareto improvement. For every $e_i \neq e$, the tax schedule near $b_i^\star$ is unchanged

(Step 2), so moving from (τ, g) to $(\hat{\tau}, \hat{g})$ affects her only through the strict rise in g. By the envelope theorem, her indirect utility is strictly increasing in g (a marginal increase raises old-age consumption one-for-one through ψg, holding her choice fixed), so she is strictly better off. For agent e: since $\hat{\tau}(b^\star) = \tau(b^\star)$, choosing $b^\star$ under $\hat{\tau}$ at the original g would exactly replicate her original utility; by revealed preference, her actual choice $\hat{b} \neq b^\star$ under $(\hat{\tau}, g)$ must be strictly better (Step 1 shows $\hat{b}$ strictly dominates $b^\star$ directly), and the further rise from g to $\hat{g}$ makes her strictly better off still, by the same envelope argument.

Thus $\hat{\tau}$ yields a strict Pareto improvement over τ, and in particular strictly higher utilitarian welfare, so τ is not optimal. $\square$

B.6 Existence of an Optimal Concave Tax Schedule

This section proves Proposition 3 by existence: I show that the set of concave schedules in $\mathcal{T}$ is compact once endowed with the topology of uniform convergence of τ and τ', and that welfare is continuous in this topology, so a maximizer exists by the extreme value theorem. Throughout, I use the norm

$$\|\tau\|_1 := \max_{b\in[0,e_n]} |\tau(b)| + \max_{b\in[0,e_n]} |\tau'(b)|,$$

which controls τ and τ' only, unlike the $\mathcal{C}_2$-norm of Definition 4, which also controls τ''. This is the relevant topology here: as the proofs of Lemmas B.1 and B.2 show, every equilibrium object in the model – each agent's optimal savings, $\mathcal{G}(\tau)$, and hence welfare – depends on τ only through τ and τ', never through τ'' directly.

Lemma B.3. *$\mathcal{T}_c := \{\tau \in \mathcal{T} : \tau$ concave$\}$ is compact in $\|\cdot\|_1$.*

Proof. Let (τ_n) be a sequence in $\mathcal{T}_c$. By Assumption 4, every τ_n satisfies $|\tau_n(b)| \leq M$ and $-1 \leq \tau_n'(b) \leq M'$ for the *same* constants M, M', independent of n (this is the uniformity across the whole class $\mathcal{T}$ that Assumption 4 builds in, not merely a bound holding for each τ_n separately), so $\{\tau_n\}$ and $\{\tau_n'\}$ are uniformly bounded. Moreover τ_n' is $\overline{M''}$-Lipschitz for $\overline{M''} := \max\{M'', (1+M')^2\} < \infty$, a constant fixed in advance by Assumption 4 and not depending on n: τ_n'' is pointwise bounded between $-(1+\tau_n')^2 \geq -(1+M')^2$ and M'' by Assumption 4's two curvature conditions. Hence $\{\tau_n\}$ and $\{\tau_n'\}$ are both uniformly bounded and equicontinuous, and the Arzelà-Ascoli theorem gives a subsequence (not relabeled) with $\tau_n \to \tau$ and $\tau_n' \to h$ uniformly, for some continuous τ, h. Uniform convergence of both τ_n and τ_n' implies τ is differentiable with $\tau' = h$, so $\tau_n \to \tau$ in $\|\cdot\|_1$.

To see $\tau \in \mathcal{T}_c$: boundedness of τ, τ' and $\tau' \geq -1$ pass to the limit since they are closed conditions under uniform convergence, and concavity (τ' non-increasing) is likewise preserved under uniform convergence of τ_n'. As a uniform limit of $\overline{M''}$-Lipschitz functions, τ' is itself $\overline{M''}$-Lipschitz, hence differentiable almost everywhere (Rademacher), and a standard difference-quotient argument (using $\tau_n' \to \tau'$ uniformly) shows the curvature bound $-(1+\tau'(b))^2 \leq \tau''(b)$ holds at every point b where $\tau''(b)$ exists, i.e., almost everywhere (concavity itself already gives $\tau'' \leq 0 \leq M''$, so the upper bound is automatic). This is all that the arguments in Appendix B that reference τ'' actually use: since τ' is Lipschitz, it is absolutely continuous, so its a.e.-derivative integrates back to τ' itself, and every conclusion drawn from the curvature bounds elsewhere in this appendix (in particular, strict global concavity of each agent's problem) goes through verbatim once stated in terms of τ' being absolutely continuous with the bound holding a.e., rather than τ'' existing at literally every point. If the resulting τ is not classically $\mathcal{C}^2$, it can be approximated arbitrarily closely, in $\|\cdot\|_1$, by a genuine $\mathcal{C}^2$, concave element of $\mathcal{T}_c$: mollify τ' with a smooth, non-negative kernel of bandwidth ρ. This yields a $\mathcal{C}^1$ function whose derivative is smooth (so the mollified τ is $\mathcal{C}^2$); it remains non-increasing, since for a non-negative kernel, convolving a non-increasing function preserves non-increasingness pointwise (for $b_1 < b_2$, $\tau'(b_1 - y) \geq \tau'(b_2 - y)$ for every y, so the same holds after averaging against the kernel); it converges to τ in $\|\cdot\|_1$ as $\rho \to 0$; and it satisfies the curvature bound up to an $O(\rho)$ correction that vanishes as $\rho \to 0$, by standard properties of mollification. This closes the gap between $\mathcal{C}^{1,1}$ and $\mathcal{C}^2$ without any loss, by continuity of welfare (Lemma B.4), so I do not distinguish between the two below. Hence $\mathcal{T}_c$ is compact. □

Lemma B.4. *Utilitarian welfare,* $W(\tau) := \sum_{i=1}^{n} \pi_i \left[u\big(c^y_{\tau,\mathcal{G}(\tau)}(e_i)\big) + \beta u\big(c^o_{\tau,\mathcal{G}(\tau)}(e_i)\big)\right]$, *is continuous on* $\mathcal{T}$ *in* $\|\cdot\|_1$.

Proof. By Lemma B.2, $\mathcal{G}(\cdot)$ and each $b_{\cdot,\mathcal{G}(\cdot)}(e_i)$ are continuous in τ under the $\mathcal{C}_2$-norm; both proofs, as written, use only the closeness of τ and τ' (the FOC characterizing $b_{\tau,g}$ depends only on τ, τ'), so the same conclusions hold under the weaker $\|\cdot\|_1$-norm. Since $b^c_{\tau,\mathcal{G}(\tau)}(e_i)$ is recovered from $b_{\tau,\mathcal{G}(\tau)}(e_i)$ and $\mathcal{G}(\tau)$ via the same continuous first-order conditions, it too is $\|\cdot\|_1$-continuous in τ. Consumption $c^y_{\tau,g}(e_i) = e_i - \tau(b_{\tau,g}(e_i)) - b_{\tau,g}(e_i) - b^c_{\tau,g}(e_i)$ and $c^o_{\tau,g}(e_i) = Rb_{\tau,g}(e_i) + b^c_{\tau,g}(e_i) + \psi g$ are continuous compositions of these objects, and u is continuous, so W is a finite sum of continuous functions of τ, hence continuous. □

Proof of Proposition 3. By Lemma B.3, $\mathcal{T}_c$ is compact; by Lemma B.4, W is continuous on $\mathcal{T}$, and hence on the subset $\mathcal{T}_c$. By the extreme value theorem, W attains a maximum on $\mathcal{T}_c$: there exists $\tau^\star \in \mathcal{T}_c$ with $W(\tau^\star) = \sup_{\mathcal{T}_c} W$. This $\tau^\star$ is concave and maximizes welfare among concave tax schedules, as required. □

B.7 Proof of Proposition 4.

1. There exists $\tilde{b}$ such that $\tau^\star(b) = \bar{\tau}$ for all $b \geq \tilde{b}$.

   Proof. I first show that if $\tau^\star$ is optimal, then it cannot be the case that every agent who chooses to have positive bank savings saves in a strictly increasing region of $\tau^\star$. That is, there must exist at least one agent who saves a strictly positive amount $b > 0$ such that $\tau^{\star\prime}(b) = 0$.

   Suppose, by way of contradiction, that this is not the case. That is, for all $e_i \in E$ such that the equilibrium bank savings satisfy $b^\star_i > 0$, it holds that $\tau^{\star\prime}(b^\star_i) > 0$. Let $\tilde{b}$ denote the largest equilibrium bank-saving level among such agents, and let $\bar{e} \in E$ be the highest-income type whose equilibrium bank savings satisfy $b^\star_{\bar{e}} = \tilde{b}$.

   Fix $\varepsilon > 0$ arbitrarily small. Consider a tax schedule $\hat{\tau} \in \mathcal{T}_c$ satisfying the following properties:

   $$\hat{\tau}(b) = \tau^\star(b) \quad \text{for all } b \leq \tilde{b}, \qquad \|\hat{\tau} - \tau^\star\| < \varepsilon, \qquad \hat{\tau}(b) < \tau^\star(b) \quad \text{for all } b > \tilde{b}.$$

   Such a construction is feasible within the set of *concave* admissible taxes $\mathcal{T}_c$, not merely within $\mathcal{T}$: since $\tau^\star$ is concave, $\tau^{\star\prime}$ is non-increasing, so for $\delta, w > 0$ small I can let $\hat{\tau}'$ agree with $\tau^{\star\prime}$ on $[0, \tilde{b}]$, taper smoothly and monotonically down from $\tau^{\star\prime}(\tilde{b})$ to $\tau^{\star\prime}(\tilde{b}) - \delta$ over $(\tilde{b}, \tilde{b} + w)$ (e.g., via a $\mathcal{C}^1$ cubic interpolation with zero derivative at both endpoints of the tapering interval, so the taper itself is non-increasing), and then continue as $\tau^{\star\prime}(b) - \delta$ on $[\tilde{b} + w, e_n]$. The taper never increases the slope, so $\hat{\tau}'$ inherits the non-increasingness of $\tau^{\star\prime}$ throughout $[0, e_n]$, making $\hat{\tau}$ concave; for δ, w small enough $\hat{\tau}$ also remains within the level and curvature bounds of Assumption 4, so $\hat{\tau} \in \mathcal{T}_c$. This matters because $\tau^\star$ is hypothesized to be optimal only *among concave schedules* (Proposition 3), so a competitor must itself lie in $\mathcal{T}_c$ for the welfare comparison below to contradict that optimality.

Since type $\bar{e}$ saves $\tilde{b} > 0$ at the optimum and $\tau^{\star\prime}(\tilde{b}) > 0$, her equilibrium choice satisfies the interior first-order condition. Lowering taxes strictly above $\tilde{b}$ relaxes this condition and induces a strictly positive marginal increase in bank savings for type $\bar{e}$. As a result, type $\bar{e}$ is strictly better off under $\hat{\tau}$.

Moreover, $\hat{\tau}(b) \leq \tau^\star(b)$ for every $b \in [0, e_n]$, with equality for $b \leq \tilde{b}$: every type faces a weakly lower tax at every possible savings level under $\hat{\tau}$ than under $\tau^\star$, holding $g = g^\star$ fixed throughout. For any other type e_i with $b_i^\star < \tilde{b}$, she can therefore always replicate her original choice $b_i^\star$ under $\hat{\tau}$ and obtain exactly her original payoff, since $\hat{\tau}(b_i^\star) = \tau^\star(b_i^\star)$. By revealed preference, her actual optimal choice under $(\hat{\tau}, g^\star)$ can only be weakly better for her. Hence no type with $b_i^\star < \tilde{b}$ is made worse off.

Therefore, aggregate welfare under $\hat{\tau}$ is strictly higher than under $\tau^\star$, contradicting the optimality of $\tau^\star$. This implies that, at an optimal tax schedule, there exists at least one agent who saves a strictly positive amount in a flat region of the tax schedule, i.e., at some $\bar{b} > 0$ with $\tau^{\star\prime}(\bar{b}) = 0$.

Since $\tau^\star$ is concave, its marginal tax rate $\tau^{\star\prime}$ is non-increasing. Hence, once $\tau^{\star\prime}(b) = 0$ at some $\bar{b}$, it must be that $\tau^{\star\prime}(b) = 0$ for all $b \geq \bar{b}$. This implies that $\tau^\star(b)$ is constant for all $b \geq \bar{b}$. Letting $\bar{b} = \tilde{b}$ completes the proof. □

2. $\tau^\star$ is non-decreasing.

   Proof. Since $\tau^\star$ is a concave tax schedule (by hypothesis, as Proposition 3 establishes that an optimal schedule within the concave class exists), $\tau^{\star\prime}$ is a non-increasing function of b. The previous part shows that $\tau^{\star\prime}(b) = 0$ for all $b \geq \tilde{b}$. Since $\tau^{\star\prime}$ is non-increasing, this implies $\tau^{\star\prime}(b) \geq \tau^{\star\prime}(\tilde{b}) = 0$ for all $b \leq \tilde{b}$ as well. Hence $\tau^{\star\prime}(b) \geq 0$ for every $b \in [0, e_n]$, i.e., $\tau^\star$ is non-decreasing. □

3. At $\tau^\star$ there is zero evasion.

   Proof. This is a direct consequence of Proposition 2. □

4. Every agent whose optimal savings are greater than or equal to $\tilde{b}$ consumes the same amount when young and old. However, this amount cannot be equal across all agents with savings bigger or equal to $\tilde{b}$.

   Proof. As in the full information case, if the tax schedule implied equal consumption across agents with different incomes, then the tax schedule would have a positive marginal tax rate (and actually taxes would be progressive), which cannot happen if $b > \tilde{b}$. □

C Numerical Implementation

C.1 Beta-Binomial Distribution

I consider the Beta-Binomial distribution $X \sim \text{BetaBin}(n, \alpha, \beta)$ which is a discretized version of the beta function, and that allows to capture a wide range of income distributions.

The Beta-Binomial distribution arises when the probability of success in a Binomial distribution is itself a random variable drawn from a Beta distribution. Let $X \sim \text{BetaBin}(n, \alpha, \beta)$. Then the probability mass function of $X \in \{0, 1, \ldots, n\}$ is given by:

$$\mathbb{P}(X = k) = \binom{n}{k} \cdot \frac{B(k + \alpha,\, n - k + \beta)}{B(\alpha,\, \beta)},$$

where $B(a, b)$ denotes the Beta function, defined in terms of the Gamma function $\Gamma(\cdot)$ as:

$$B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a + b)}.$$

Using this expression, the Beta-Binomial probability mass function can also be written as:

$$\mathbb{P}(X = k) = \binom{n}{k} \cdot \frac{\Gamma(k + \alpha)\Gamma(n - k + \beta)\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(n + \alpha + \beta)}.$$

The Gamma function $\Gamma(z)$ is a continuous extension of the factorial function, defined for all $z > 0$ by:

$$\Gamma(z) = \int_0^\infty t^{z-1} e^{-t}\, dt.$$

It satisfies $\Gamma(n) = (n - 1)!$ for all positive integers n.

If $X \sim \text{BetaBin}(n, \alpha, \beta)$ and $\alpha > \beta$ then the distribution will be left-skewed (have higher probability mass in higher incomes), while if $\alpha < \beta$ then the PDF is right-skewed. An example is illustrated in Figure XII.

C.2 Numerical Implementation

To numerically solve the government's problem, I proceed in several steps, addressing both the individual optimization problems and the overall variational structure of the planner's objective.

First, for a given tax schedule τ and public good provision level g, I solve each agent's problem by considering a discrete grid over (b, b^c), where b denotes formal savings and b^c denotes informal (couch) savings. At each grid point, I compute the agent's utility and determine the optimal allocation by maximizing the objective function subject to the feasibility constraints. This yields, for every agent type $e \in E$, an optimal pair $(b_{\tau,g}(e), b^c_{\tau,g}(e))$.

Using these optimal savings decisions, I then compute the implied public good contribution as a function of τ and g, which I denote by $\text{PGC}(\tau, g) = \sum_{e \in E} \tau(b_{\tau,g}(e))\pi(e)$. An equilibrium level of the public good must satisfy the consistency condition $g = \text{PGC}(\tau, g)$, so I compute g^* as a fixed point of the mapping $g \mapsto \text{PGC}(\tau, g)$. This is done using MATLAB's `fsolve` (default

Figure XII: Beta-Binomial Density Function.

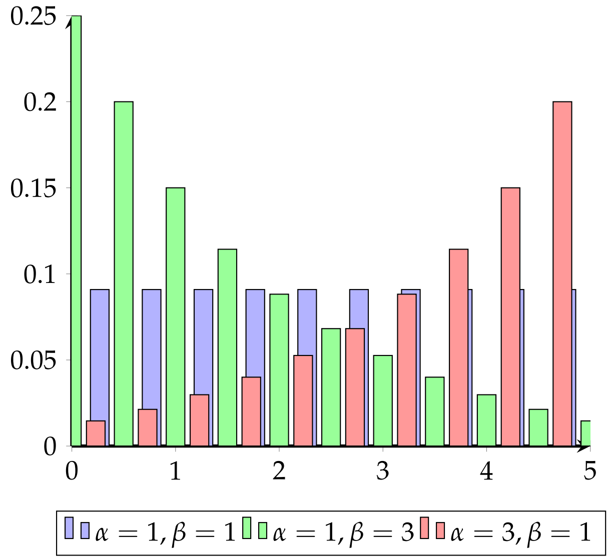

function and step tolerances of 10^{-6}) applied to $\mathrm{PGC}(\tau, g) - g$, starting from an initial guess of $g = 0.5$ in every case.

The government's problem, in turn, is a variational problem: it involves maximizing a functional over a space of admissible tax schedules. To solve this numerically, I implement Ritz's method (see Gelfand and Fomin (1963)), which approximates the infinite-dimensional optimization by projecting the tax function τ onto a finite-dimensional space spanned by basis functions. In particular, I approximate τ using Chebyshev polynomials of the first kind up to order five. This choice ensures a good balance between flexibility and numerical tractability, and it also facilitates enforcing smoothness constraints on the derivatives of τ.

The optimization is subject to several constraints. First, the incentive-compatibility conditions of agents must hold, meaning that their savings decisions are indeed optimal given τ and the induced g; this is enforced by construction, since g is always recomputed as the fixed point of $\mathrm{PGC}(\tau, \cdot)$ and each agent's (b, b^c) is always recomputed as the grid maximizer described above whenever τ is updated during the search. Second, τ must lie in the set $\mathcal{T}$, which imposes bounds on both the level and the first and second derivatives of the tax schedule; the specific bounds used in the search are described below, together with the parameterization to which they are applied.

The Chebyshev-basis approximation described above is what I use for the unrestricted numerical check discussed after Proposition 3 (footnote 3), where I verify that an unconstrained search over this basis returns a concave schedule in every parameterization considered. For the constrained, flat-tail characterization that generates the figures and the table reported in Section 6, I instead use a lower-dimensional, hybrid parameterization that directly imposes the qualitative shape of Proposition 4, described next.

Baseline calibration. Unless noted otherwise in a figure or table, the results reported in Section 6 use the following common specification:

- *Utility:* CRRA utility, $u(c) = (c^{1-\sigma} - 1)/(1-\sigma)$ (with the $\log(c)$ limit at $\sigma = 1$), with $\sigma = 6$. A candidate allocation with negative consumption in either period is assigned a large negative utility penalty rather than excluded from the search directly, so infeasible choices are simply dominated rather than causing the solver to fail. I flag as a caveat that, because $\sigma = 6$ exceeds the income levels considered below, Assumption 3's bound $0 < A(c) < 1$ is not satisfied pointwise everywhere along the resulting consumption paths; nothing in the numerical procedure itself relies on Assumption 3, so the qualitative results are unaffected, but I note the tension for transparency.

- *Returns:* $\beta = 1/R$ (Assumption 1), and $\psi = 1.2R$, so that $\psi > R$ holds with a fixed 20% margin throughout.

- *Income:* $n = 20$ equally spaced income types on $E \in [1, 5]$, with probabilities $\{\pi_i\}$ given by the discretized Beta-Binomial distribution $\text{BetaBin}(n-1, \alpha, \beta_d)$ described above.

- *Savings grid:* each agent's problem is solved over a grid of 200 equally spaced points for b on $[0, e_n]$ and 40 equally spaced points for b^c on $[0, e_n]$ (a coarser grid for b^c, since couch savings play a secondary role once the tax schedule is admissible).

Tax parameterization. For $b \le \tilde{b}$, I parameterize τ as a quadratic function, $\tau(b) = ab^2 + cb + d$; for $b > \tilde{b}$, $\tau(b) = \bar{\tau}$ is constant. Rather than treating $a, c, d, \tilde{b}$, and $\bar{\tau}$ as five free parameters, I impose the smooth-pasting conditions $\tau(\tilde{b}) = \bar{\tau}$ and $\tau'(\tilde{b}) = 0$ directly, which pin down $c = -2a\tilde{b}$ and $d = \bar{\tau} - c\tilde{b} - a\tilde{b}^2$ as functions of $(a, \tilde{b}, \bar{\tau})$; the search is then over these three free parameters only. This construction matches τ and τ' continuously at $\tilde{b}$, but not τ'': the quadratic branch has $\tau''(b) = 2a$ throughout, while the constant branch has $\tau''(b) = 0$, so τ'' is generally discontinuous at $\tilde{b}$ unless $a = 0$. The resulting numerical schedule is therefore $\mathcal{C}^{1,1}$ rather than strictly $\mathcal{C}^2$ at the junction point, the same gap discussed in Section B.6 for the theoretical existence argument.

The search over $(a, \tilde{b}, \bar{\tau})$ is subject to the bound constraints $a \in [-1/2, 0]$, $\tilde{b} \in [0, e_n]$, and $\bar{\tau} \in [0, e_n]$, together with nonlinear constraints enforcing the curvature condition of Assumption 4 on the quadratic branch and $\tau(0) \ge 0$ (ruling out the negative-revenue schedules discussed in Section B.2). Because $a \le 0$ is already imposed as a bound, every schedule in this family is automatically concave with $\tau'' \le 0$, so I do not separately calibrate the upper curvature bound M'' of Assumption 4: any $M'' \ge 0$ is satisfied without further restriction.

I solve for the welfare-maximizing $(a, \tilde{b}, \bar{\tau})$ using MATLAB's `fmincon` (SQP algorithm, capped at 100 iterations and 1,000 function evaluations), wrapped in a `MultiStart` search over 3 randomly drawn feasible starting points to guard against convergence to a poor local optimum. A candidate solution is retained only if the solver reports a valid exit flag and a first-order optimality measure between 10^{-11} and 1.5; each optimal-tax figure and each row of Table I reports the specification that passes this check for the corresponding (R, α, β_d) combination. Table I itself is constructed by re-running this entire procedure for each pairing of $R \in \{1.05, 1.15\}$ and $(\alpha, \beta_d) \in \{(1,1), (1,3), (3,1)\}$ and recording the resulting $g^\star$ and financial informality rate.

Step-tax exercise. The comparison in Section 5.4 uses the same solution procedure (agent-level grid search, budget-balancing `fsolve`, and `fmincon/MultiStart` welfare maximization) applied instead to a genuinely piecewise-constant schedule with $N \in \{2, 3, 5\}$ steps: N tax levels and $N - 1$ breakpoints are chosen jointly to maximize utilitarian welfare, subject to the same budget-balance and grid-search machinery described above, with no smoothness restriction imposed between steps (such a schedule is discontinuous and hence outside $\mathcal{T}$). Figure VII reports the $N = 5$ case.