

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

Alberto Ramirez de Aguilar[†]

January 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 informality and tax evasion arise endogenously as responses to the tax schedule, and I characterize the welfare-maximizing tax structure under these considerations. The optimal tax is concave and 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. The model suggests that the optimal policy is to rationalize and tolerate some level of 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 with large informal sectors.

*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.

[†]Banco de México. E-mail: aramirez@banxico.org.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 with large informal sectors, where individuals often operate outside the formal system and income-based taxation becomes infeasible. In such settings, governments must rely on observable proxies, such as deposits in the banking system, to collect taxes. This creates an incentive for agents to conceal income through informal 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, informality and evasion arise endogenously: agents may avoid the financial system altogether or use it strategically to reduce their tax burden. 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 interest rate than the couch. Income when old is comprised by 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, having as an objective to maximize utilitarian welfare. The optimal tax is progressive: richer agents face higher marginal tax rates, enabling redistribution without distorting savings. Since the taxed bank account earns interest while the untaxed account does not, agents always prefer to save formally, and 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 informality and evasion, which limits the government's ability to redistribute. I show that the optimal tax is concave and flat beyond a threshold. This structure keeps richer agents in the formal sector, avoids evasion at the top, and maintains participation at the bottom. Even though the government values both redistribution and revenue (through public good provision), it cannot impose progressive taxes without inducing costly evasion. The outside option thus shapes the optimal policy and limits the feasibility of progressive taxation.

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 informal savings and endogenous evasion. Second, it offers a theoretical complement to empirical work on informality, highlighting how tax policy can shape sectoral choices. Third, it provides normative guidance for developing economies, where large informal sectors and limited enforcement capacity make the trade-offs between redistribution, tax collection, and public goods particularly acute.

This paper fits within the extensive literature on optimal taxation and asymmetric information. Classical taxation models such as [Mirrlees \(1971\)](#) and [Diamond and Mirrlees \(1971\)](#) analyze optimal income taxation, assuming full income observability and an impediment for the government to observe individual labor choices. 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). The classic result of this early literature is that the optimal tax has zero marginal tax rates at the top (i.e., eventually the tax becomes constant), since progressive taxes (ones with increasing marginal rates) would incentivize high productivity individuals to lower their effort choices, and hence their income, which would result in both a gain due to be working less and paying fewer taxes.

In classic Mirrleesian-type economies, providing general statements about the optimal tax schedule's overall shape is difficult since taxes are usually U-shaped. 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, resulting in concave taxes.

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 faces uncertainty 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 is concave. This is because agents in my model have an outside option (informal savings) that they can turn to if taxes become too progressive. This outside option constrains the tax authority, making progressive taxes counterproductive as they induce evasion and reduce public good provision. 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.

Several other papers have extended the Mirrlees framework to study optimal taxation under various frictions and informational constraints. [Albanesi and Sleet \(2006\)](#) analyze dynamic optimal taxation with private information and emphasize the role of intertemporal distortions in shaping long-run redistribution. Recent work by [Golosov and Tsyvinski \(2006\)](#) and [Kocherlakota \(2010\)](#) studies optimal taxation in dynamic economies with private information. These models focus on insurance mechanisms and intertemporal distortions arising from tax design. [Farhi and Gabaix \(2020\)](#) introduce behavioral agents into the Mirrlees setting and show how cognitive frictions affect the structure of optimal taxes. Like [Vairo \(2025\)](#), these papers explore settings where the planner departs from purely utilitarian objectives or faces non-standard constraints, leading to tax schedules that may differ from the classic concave benchmark. My paper contributes to this literature by introducing a new friction—“hidden savings”—that interacts with the government’s redistributive and revenue goals in novel ways, leading to new insights about the feasibility and desirability of progressive taxation.

The main application of my framework concerns how the tax system affects the formation of the informal sector and whether or not zero informality is optimal from a utilitarian welfare perspective. Defining what informality means is not an easy problem at all. According to surveys and reports of the World Bank, like [Perry et al. \(2007\)](#) and [Oviedo and Thomas \(2009\)](#), informality can be defined as “any economic activity taking place below the radar of the government”. This definition is very broad, so, as these reports clarify, there is usually two different ways to further define informality: from a point of view of firms, referred to as the “productivity view”, and from the point of view of workers, referred to as the “social protection view”. The productivity view usually focuses on different types of firms and their legal status. In contrast, the social protection view focuses on whether or not workers have access to social security, pay taxes according to their income, and have access to all the compensations they should receive according to the law. In this paper, I present a reduced-form model of informality that does not consider firms. I consider an agent informal whenever it uses her couch (cash) to hide its income and savings from the tax authority.

But what causes informality? According to [Ulyssea \(2020\)](#), until not so long ago, informality was seen as a result of a worker being excluded from the formal sector, and hence, becoming informal was a residual choice between working and receiving a low payment or being left unemployed. However, recent work has shown that this hypothesis does not capture all the forces pushing towards informality. The most accepted thesis regarding the cause of informality nowadays is that informality is caused by burdensome regulations (e.g., high entry costs, high taxes, complicated enrollment procedures) to become formal. Hence, workers choose to be informal to avoid such regulations. In this sense, [Oviedo and Thomas \(2009\)](#) shows empirical evidence suggesting that many informal workers choose not to be formal since, if they work in the formal sector, they would earn a similar wage but would have to pay taxes, reducing their income. Following this line of thought, [Bosch and Maloney \(2008\)](#) describes both sectors not as substitutes but as complements and shows, using data from Brazil and Mexico, that there is a considerable turnout of workers between these sectors. In my paper, this sector-choice decision that agents have, which is captured by the type of saving system the worker chooses (cash or in the bank), is going to be a response of the tax system the agent faces, and the amount of taxes it would imply to pay if the worker chooses to save in the bank.

Another problem with informality is the free-rider problem of public goods provision. A large informal sector means that a significant proportion of the population is either not paying or

underpaying the taxes they are supposed to pay. Hence, overall tax collection decreases as the informal sector grows. This has consequences regarding public goods provision since they can be used by any individual regardless of whether they pay taxes. In this sense, if the government could somehow tax informal workers, this would help mitigate the free-rider problem. However, taxing informal workers is not easy since, in many cases, the government does not even know the income of informal agents, as [Meghir et al. \(2015\)](#) suggest. If it is not done cautiously, changing the tax system could lead to an even higher informality rate. Nevertheless, there is empirical evidence on how some local governments have dealt with the taxation problem that informality brings. In particular, [Olken and Singhal \(2011\)](#) present evidence using survey micro-data on how, through non-conventional mechanisms, local governments have been able to collect money (taxes) from informal workers and that this increases public revenue and the tax burden for formal workers. The authors show that these types of “informal taxes” tend to be regressive for informal workers (i.e., the quotient tax/income decreases with income), but nonetheless, people do pay them. In my model, the tax burden imposed to low-income agents crucially depends on the income distribution and its support. As there are wealthier individuals in the economy, the government can collect taxes from richer agents and reduce the tax burden on low-income workers.

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.

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. Agents receive an income when born, independent, and identically distributed from other agents. Income has a finite support $E = \{e_1, \dots, e_n\}$ with $e_1 < e_2 < \dots < e_n$. The probability that agent $i \in [0, 1]$ has an income e_j is given by $0 < \pi_j < 1$. I denote F to be the distribution associated with the type assignment.

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 rate of $R > 1$; or they can save in their “couch,” which pays an interest rate of one. 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 .

One particular type of tax that may sound appealing from a “social justice” point of view 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 $e \in \{e_2, \dots, e_n\}$ 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 informality/evasion and to see if the welfare-maximizing tax implies no informality (and hence from a normative point of view, informality could be thought of as “bad” for the economy’s welfare). In the context of my model, I define an agent to be **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 conceiving her income from the government (which is why I call this agent informal). In contrast, if the agent does save but does not completely “report” her entire savings, the agent is an evader.

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 $\tau : E \rightarrow \mathbb{R}$ 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_j is:

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

where $b, b^c \geq 0$ represent the amount of bank/couch savings; $\tau(e_j)$ are the taxes that an individual of type e_j 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 $R \leq \psi$, which guarantees that a positive amount of public good is provided in equilibrium, and therefore, collecting taxes is a non-trivial problem.

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 $\tau : E \rightarrow \mathbb{R}$ 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 sounds 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:

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

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

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

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

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,

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

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.

Lemma 1. *Let $\tau : E \rightarrow \mathbb{R}$ 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).*

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 $E \sim U([1, 2])$.

The proof of [Lemma 1](#) can be found in [Section 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:

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

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

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

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 [Section A](#).

Proposition 1. *The optimal tax schedule τ^* has the following properties.*

1. *It is unique.*
2. *It equates consumption when young across all agents.*
3. *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.*
4. *If all agents have positive savings given τ^* , then the optimal tax schedule τ^* is progressive.*
5. *It implies zero informality and zero evasion.*

This proposition states that the optimal tax schedule τ^* implies the same consumption across different agents in the economy, which means that the optimal taxes should promote perfect equality in the economy. The second 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 .¹


The last statement of this proposition says that the optimal tax schedule induces agents not to be 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.

One parameter that I want to highlight is the interest rate R and how changes in it affect the optimal tax. As R increases, saving becomes more attractive since it has higher returns when old. On the other hand, since savings are higher, in order to equate consumption across periods, the government cannot charge high taxes, since otherwise consumption when young would be low. Hence, as the interest rate increases, we should expect a decrease in the overall tax collection, which translates into a loss in public good provision. When R increases, agents are then not using that much the public good to hedge old consumption; they use more of their own savings. This can be seen in [Figure II](#), which shows that as R increases, the optimal tax schedule decreases, and both savings and consumption (when old and young, which are equal) increase.

As the interest rate increases, since private savings provide a higher interest rate payment when old, there is less need to collect higher overall taxes and to provide a large amount of the public good. As a result, the optimal tax collection $\sum_{j=1}^n \tau(e_j)\pi_j$ is a decreasing function of R . This can be seen in [Figure III](#).


¹ However, the optimal value of consumption across agents does depend on $\{\pi_j\}_j$.

Figure II: OPTIMAL TAXES AND THE INTEREST RATE.



NOTES: This graph plots the optimal tax schedule for two different values of the interest rate: a low interest rate $R_L = 1.01$ displayed in blue, and a high interest rate $R_H = 1.1$ displayed in green. The second panel plots each agent savings in the bank, while the third panel plots consumption (when old and young, which are equal). For this figure, I consider $\psi = 1.2$, and $E \sim U([1, 2])$.

Figure III: TAXES AND PUBLIC GOOD AS A FUNCTION OF THE INTEREST RATE.



NOTES: The left panel of this figure plots the optimal good provision as a function of the interest rate R , while the right panel displays the optimal tax collection g as a function of R . For this figure, I consider $\psi = 1.2$, and $E \sim U([0, 1])$.

5 Hidden Income 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 $\tau : [0, e_n] \rightarrow \mathbb{R}$, where the domain represents the bank savings of each agent. Agents cannot borrow and the maximum amount an agent could (potentially) save is e_n .

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 where the informality rate is very high, 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;² second, nature assigns each agent an income type according to the distribution F ; and second, 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:

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

where b, b^c are restricted to be non-negative. 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 informality and evasion.

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

The idea behind these definitions is that agents who do not use the financial system at all and rely on the couch as a means of saving are those who seek to completely conceal their income from the government, thereby making them 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.

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.³ As shown in [Figure II](#), in the benchmark framework, the optimal tax schedule τ is such that every agent with positive savings has the same amount of money in the bank, \hat{b} . Also, let \hat{g} be the amount of public good provided under τ . If this same tax schedule were to be implementable under hidden information, then $\hat{g} = \tau(\hat{b})$, since all agents have the same savings level. If τ is increasing at \hat{b} , since from the point of view of each agent \hat{g} is taken as given, then at least one agent (most likely an agent with low income) has an incentive to reduce their savings a little bit, which almost does not alter their income when old, but decreases taxes, increasing consumption when young. In a similar way, if τ is decreasing at τ then the richer agents have an incentive to save more, and both increase their young and old consumption. On top of this, there is an additional complication for the

2 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 \(1987\)](#).

3 [Section 5.4](#) argues that non-continuous taxes are also not optimal.

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.

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:

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

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 and that the utility function is steeper. The following assumption imposes an upper and lower bound on the values that $A(c)$ can take.

Assumption 3. The utility function $u : \mathcal{D} \rightarrow \mathbb{R}$ 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.⁴

⁴ This follows from the fact that for CRRA utility, $A(c) = \sigma/c$.

Assumption 4. I consider functions $\tau : [0, e_n] \rightarrow \mathbb{R}$ that are continuously twice differentiable, bounded, have uniformly bounded first derivatives, and:

$$\begin{aligned}\tau'(b) &\geq -1 \quad \text{for all } b \in [0, e_n], \\ \tau''(b) + (1 + \tau'(b))^2 &\geq 0 \quad \text{for all } b \in [0, e_n].\end{aligned}$$

I denote \mathcal{T} as the set of functions satisfying these conditions. 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.⁵ 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 is strictly concave and differentiable. Section 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. The main restriction that \mathcal{T} imposes is to discard functions whose concavity (and hence $\tau''(b) < 0$) reverses the strict concavity of an agent's optimization problem.

5.1 Equilibrium and Government's Problem

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

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

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.

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

1. Given $\tau, g; (b_{\tau,g}(e), b_{\tau,g}^c(e))$ solve:

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

for every $e \in E$.

⁵ Let $f \in \mathcal{C}_2([a, b])$ and define $\|\cdot\| : \mathcal{C}_2([a, b]) \rightarrow \mathbb{R}_+$ 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.

2. There is budget balance:

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

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

Section B shows that if $\tau \in \mathcal{T}$ then $\mathcal{G}(\tau) \neq \emptyset$. Furthermore, the mapping $\mathcal{G} : \mathcal{T} \rightarrow \mathbb{R}_+$ is actually a function (i.e., each tax function is assigned exactly to one g) and it is continuous.⁶

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:

$$\begin{aligned} \max_{\tau \in \mathcal{T}} \sum_{i=1}^n \left[u \left(e_i - \tau(b_{\tau,g}(e_i)) - b_{\tau,g}(e_i) - b_{\tau,g}^c(e_i) \right) + \beta u \left(Rb_{\tau,g}(e_i) + b_{\tau,g}^c(e_i) + \psi g \right) \right] \pi_i, \quad \text{s.t.:} \\ (b_{\tau,g}(e_i), b_{\tau,g}^c(e_i)) \in \operatorname{argmax}_{b, b^c \geq 0} u(e_i - b - b^c - \tau(b)) + \beta u(Rb + b^c + \psi g) \quad \text{for every } e_i \in E, \\ g = \mathcal{G}(\tau). \end{aligned}$$

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,⁷ and since the objective function is continuous in τ ,⁸ 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.

⁶ This function is continuous with respect to the \mathcal{C}_2 -norm defined in the twice-continuously differentiable functions. For more details consult [Section B](#).

⁷ Restricting each tax function to have a compact domain together with each function having a uniformly bounded first derivative, implies compactness via the Arzelá-Ascoli Theorem. More details can be found in ??.

⁸ 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 [Section B](#).

Lemma 2. Let $\tau \in \mathcal{T}$, $g > 0$, and $(b_{\tau,g}(e), b_{\tau,g}^c(e))$ 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 $b_{\tau,g}^c(e) = 0$ then:

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

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

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

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

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

The proof of this lemma can be found in [Section 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 $(b_{\tau,g}(e), b_{\tau,g}^c(e))$ 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. $b_{\tau,g}^c(\cdot)$ is a non-increasing function of g .

[Section B](#) provides the proof of this lemma. This lemma states that bank savings cannot decrease with income, suggesting that richer individuals should save at least the same amount as poorer agents in both the couch and 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. 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 $\psi \geq R$, 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 $(b_{\tau,g}(e), b_{\tau,g}^c(e))$ be the solution of income type e 's problem. Let $c_{\tau,g}^y(e) = e - \tau(b_{\tau,g}(e)) - b_{\tau,g}(e) - b_{\tau,g}^c(e)$ and $c_{\tau,g}^o(e) = Rb_{\tau,g}(e) + b_{\tau,g}^c(e) + \psi g$ be the optimal consumption when young and old (respectively) of agent $e \in E$.

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

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 [Section 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 informal savings. This section characterizes how such a policy can be optimally designed.⁹

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 cannot be optimal.

Proposition 2. Let $\tau \in \mathcal{T}$ be such that there is at least one $e \in E$ with $b_{\tau,g}(e) > 0$, $b_{\tau,g}^c(e) > 0$, 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

⁹ A complete, closed-form characterization of the optimal tax schedule under affine taxation is provided in ?? whose insights extend to the general case is provided in ??.

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 couch for bank 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 $\hat{\tau}$. [Section 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.¹⁰

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 informality). Whether low-income individuals become 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 sector of the economy. 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 I want 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 IV](#) 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.

Now, I describe the shape of the optimal tax schedule. In this section I provide the characterization of the optimal tax schedule, while in the following section I describe how the shape of the optimal tax schedule changes as I vary the model's primitives. I also discuss in [Section 6](#) the implications for the informal/formal sector and my model's lessons for this discussion.

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 V](#) 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.

¹⁰ In addition, ?? provides an explicit construction (closed form) of the welfare-improving tax schedule for the case of affine taxes.

Figure IV: OPTIMAL TAX SCHEDULE IS NOT UNIQUE.




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 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.

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 VI 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 in amount 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.


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, which is why the optimal tax is not progressive (a formal argument is

Figure V: OPTIMAL TAXES CANNOT BE DECREASING.



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 $E \sim U([1, 5])$.

Figure VI: OPTIMAL TAXES CANNOT BE PROGRESSIVE.



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 $E \sim U([1, 5])$.

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).

The following proposition shows that, without loss of generality, I can focus the analysis on concave tax schedules.

Proposition 3. *For any optimal $\tilde{\tau}^*$ there exists $\tau^* \in \mathcal{T}$ that is concave which induces the same equilibrium outcomes as $\tilde{\tau}^*$.*

The logic behind [Proposition 3](#) is that, although agents' saving choices are characterized by local first-order conditions, the construction of a concave tax schedule must guarantee that no agent finds it profitable to deviate *globally* to a different saving level. Consider an optimal tax schedule and the equilibrium saving choices it induces. For each agent, optimality requires that the marginal tax rate at her chosen saving level makes deviations in either direction unattractive: moving to higher savings would increase taxes too much relative to the return on bank savings, while moving to lower savings would reduce intertemporal smoothing. Because preferences are strictly concave and the return to bank savings is linear, each agent's equilibrium saving choice is a strict global optimum.

The key step in the construction is to modify the tax schedule only in regions where no agent is locally indifferent between alternative saving levels. In particular, non-concavities in the original tax schedule necessarily correspond to regions in which the marginal tax rate evolves in a way that weakens incentives (either by exceeding the net return on bank savings or by creating non-monotonic marginal incentives across saving levels). Replacing these regions with their concave envelope weakly relaxes marginal incentives, but never in a way that creates new profitable deviations to distant savings levels. Intuitively, the concave modification smooths the tax schedule without introducing low-tax regions that agents could exploit by reallocating their savings.

Moreover, since the modified tax schedule coincides with the original one in both levels and slopes in neighborhoods of all equilibrium saving choices, any potential gains from moving to distant regions are dominated by the loss from deviating from the locally optimal saving level. As a result, each agent's original saving choice remains globally optimal under the new tax schedule. Because equilibrium portfolios are unchanged, aggregate tax revenue and the induced level of public good provision are also unchanged.

The economic intuition is that non-concave features of the tax schedule generate unnecessary distortions in a model with hidden savings: they either encourage evasion toward the untaxed instrument or create excessive incentives to reshuffle savings across the tax base. Concavity disciplines marginal incentives while preserving implementation, allowing the government to restrict attention to concave tax schedules without loss of optimality.


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

Proposition 4. Let τ^* be an optimal concave tax schedule. Then:

1. τ^* is non-decreasing.
2. There exists \tilde{b} such that $\tau^*(b) = \bar{\tau}$ for all $b \geq \tilde{b}$.
3. At τ^* 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 optimal tax must be non-increasing since this implies that bank savings are increasing with respect to income. Then richer agents save more and pay more taxes, which is a way of redistributing income across agents. 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, which also helps to increase total utility since $\beta < 1$. Now, 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).

Figure VII: OPTIMAL CONCAVE TAX SCHEDULE $\tau^*(b)$.



The fact τ^* becomes constant for large savings comes 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 optimal income taxes in the presence of incomplete information about the worker’s ability and labor choices must have the property of zero marginal tax rates at the top. Intuitively, this happens since constant taxes incentivize high-productivity agents to work more and, hence, have higher incomes. 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.

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 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 informal agents. This section argues that the concave shape between $\tau^*(0)$ 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 $\bar{\tau}$ may be able to implement the same allocation as τ^* .¹¹ But, as [Figure VIII](#) suggests, this is not the case. The main problem with this tax schedule is that it creates evasion: agents who under τ^* 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 $\tau^*(0)$ 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 τ^* , 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 VIII](#), 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, it turns out to converge to the optimal tax schedule characterized by [Proposition 4](#) as the number of allowed steps increases. [Figure IX](#) presents a numerical example of this, when I searched for the optimal five-step tax function. As the number of allowed steps increases, the purple tax approximates the blue tax.

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

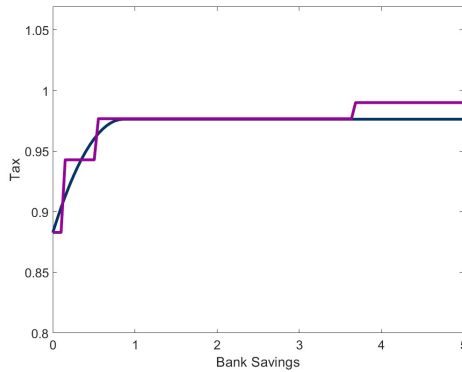
¹¹ This tax function is not continuous nor differentiable, and hence it is not an element of \mathcal{T} . However, since taxes are often step functions in many economies, I consider this class of taxes important and discuss why they are not optimal in the presence of hidden savings.

Figure VIII: 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 Section C.

Figure IX: 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 Section 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 [Section 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 informality rate. If an agent has positive couch savings but zero bank savings, I interpret this agent as 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 X](#) 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 \bar{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 $\tau^*(0) < \bar{\tau}$ 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 $b = 0$ ($\tau'^*(0) > R - 1$) 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 τ^* concavely increases up to $\bar{\tau}$ ensures that agents that have bank savings slightly higher than \bar{b} actually want to remain saving in the bank and not become evaders.¹²


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 X](#), 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 XI](#). 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 cannot allow itself to collect taxes from low-income agents. Again, the shape of the tax schedule for $0 < b < \bar{b}$ 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 individuals can advo-


¹² [Section 5.4](#) discusses this in more detail.

Figure X: 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 XI: 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 Section C.

cate collecting taxes from wealthier individuals and burden very low-income individuals with taxes.

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 Section C.

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

Income Distribution	Description	Interest Rate	g^*	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 taxes the government can charge and incentivize agents to save in the bank. If marginal taxes (and the overall tax level) 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 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 informality is not generally optimal*, even for a welfare-maximizing government. In economies where income is unobservable and savings can be hidden, 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 informal. From a normative perspective, this implies that policies aimed at fully eradicating informality may be welfare-reducing when they ignore agents' ϵ^{TM} outside options and the distortionary effects of taxing observable financial activity.

A second implication is that the *size of the informal sector depends critically on the distribution of wealth*. The quantitative results show that economies with a larger mass of high-income agents can sustain higher public good provision while tolerating lower informality rates, as the government can rely more heavily on taxing richer households without inducing widespread exit from the formal sector. 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 informality. This suggests that cross-country differences in informality need not reflect differences in enforcement capacity or institutional quality alone, but may instead be tightly linked to underlying distributional characteristics.

More broadly, the model implies that *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 informality at the bottom, the government relaxes

participation constraints and avoids inducing inefficient evasion at the top. In this sense, 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 informality. Policies that increase the relative attractiveness of formal saving—such as improving financial access or reducing intermediation costs—can therefore expand the feasibility set.

7 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 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, progressive taxation becomes counterproductive, incentivizing wealthier individuals to avoid the formal financial system, limiting revenue and redistribution potential.

The optimal tax structure that emerges is concave and flat beyond a threshold, a form that simultaneously deters evasion at the top and limits informality at the bottom. This policy design ensures that high-income individuals continue participating in the formal sector 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 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 developing countries with large informal sectors and limited enforcement capacity. The model suggests that governments can enhance welfare not necessarily by eliminating 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 general equilibrium effects, political economy constraints, or richer preferences and enforcement mechanisms heterogeneity.

References

- Albanesi, Stefania and Christopher Sleet**, "Dynamic Optimal Taxation with Private Information," *Review of Economic Studies*, 2006, 73 (1), 1–30.
- Bosch, Mariano and William Maloney**, "Cyclical movements in unemployment and informality in developing countries," Policy Research Working Paper Series 4648, The World Bank June 2008.

- Diamond, Peter A and James A Mirrlees**, “Optimal Taxation and Public Production II: Tax Rules,” *American Economic Review*, June 1971, 61 (3), 261–278.
- 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 and Aleh Tsyvinski**, “Designing Optimal Tax Systems: A Dynamic Programming Approach,” *Review of Economic Studies*, 2006, 73 (2), 365–392.
- Kocherlakota, Narayana R.**, “The New Dynamic Public Finance: A User’s Guide,” *Journal of Economic Literature*, 2010, 48 (3), 773–810.
- Meghir, Costas, Renata Narita, and Jean-Marc Robin**, “Wages and Informality in Developing Countries,” *American Economic Review*, April 2015, 105 (4), 1509–46.
- 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.
- Olken, Benjamin A. and Monica Singhal**, “Informal Taxation,” *American Economic Journal: Applied Economics*, October 2011, 3 (4), 1–28.
- Oviedo, Ana Maria and Mark Thomas**, “Economic informality: Causes, costs, and policies a literature survey,” *World Bank Working Papers*, 2009.
- Perry, Guillermo E., William F. Maloney, Omar S. Arias, and Pablo Fajnzylber**, *Informality : Exit and Exclusion* number 6730. In ‘World Bank Publications - Books.’, The World Bank Group, December 2007.
- Ramirez de Aguilar, Alberto**, “Debt, Inflation, and Government Reputation,” Working Paper, University of Pennsylvania 2025.
- Rogerson, William P.**, “The First-Order Approach to Principal-Agent Problems,” *Econometrica*, November 1985, 53 (6), 1357–1367.
- Stiglitz, Joseph E.**, “The Causes and Consequences of the Dependence of Quality on Price,” *Journal of Economic Literature*, 1987, 25 (1), 1–48.
- Ulyssea, Gabriel**, “Informality: Causes and Consequences for Development,” *Annual Review of Economics*, August 2020, 12 (1), 525–546.
- Vairo, Maren**, “Robustly Optimal Income Taxation,” *Working Paper*, 2025.

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 $h_1(b, b^c) = e - \tau(e) - b - b^c, h_2 = Rb + b^c + \psi g$ 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 $R > 1$.¹³ 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 's problem (given the tax schedule) is:

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

and this implies the following first-order conditions:

$$-u'(e - \tau(e) - b(e_k)) + u'(Rb(e_k) + \psi g) + \lambda_1 = 0, \quad (\text{A1})$$

$$-u'(e - \tau(e) - b(e_k)) + \beta u'(Rb(e_k) + \psi g) + \lambda_2 = 0, \quad (\text{A2})$$

$$b \geq 0, \lambda_1 \geq 0, \lambda_1 b = 0, \quad (\text{A3})$$

$$b^c \geq 0, \lambda_2 \geq 0, \lambda_2 b^c = 0. \quad (\text{A4})$$

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:

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

$$u'(e - \tau(e) - b(e_k)) = \beta u'(Rb(e_k) + \psi g),$$

which then implies $\beta = 1$, which 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:

$$u'(e - \tau(e) - b(e_k)) = \beta u'(Rb(e_k) + \psi g),$$

which then follows $u'(Rb(e_k) + \psi g) < u'(e - \tau(e) - b(e_k))$ since $\beta < 1$. But then, [Equation \(A1\)](#) implies that:

$$\lambda_1 = u'(e - \tau(e) - b(e_k)) - u'(Rb(e_k) + \psi g) < 0,$$

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 .

¹³ For more details, consult Appendix A of [Ramirez de Aguilar \(2025\)](#).

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

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

Since u is strictly increasing, and therefore injective, $e - \tau(e) - b(e_k) = Rb(e_k) + \psi g$. Solving for b we get:

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

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

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

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

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

And since u is strictly increasing, this 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:

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

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

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

If we substitute these constraints into the objective function, the problem becomes optimizing a strictly concave function that is evaluated at linear functions of the unknowns $\{\tau(e_j)\}_j$. Hence, as a function of these unknowns, the problem is also strictly concave, which implies that the solution is unique. □

2. The optimal tax schedule τ^* equates consumption across all agents with positive savings.

Proof. The problem I seek to solve is:

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

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

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

Let $\mathcal{B} = \{j | b(e_j) > 0\}$. Then, using the expression for $b(e_j)$, we get that if $j \in \mathcal{B}$:

$$c_j^y = c_j^o = c_j = \frac{Re_j - \tau(e_j) + \psi R \sum_{i=1}^n \tau(e_i) \pi_i}{1 + R},$$

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

$$c_j^y = e_j - \tau(e_j), \quad c_j^o = R\psi \sum_{i=1}^n \tau(e_i) \pi_i.$$

Substituting this, then the objective function becomes:

$$\begin{aligned} & \sum_{j \notin \mathcal{B}} \left[u(e_j - \tau(e_j)) + \beta u \left(R\psi \sum_{i=1}^n \tau(e_i) \pi_i \right) \right] \pi_j + \\ & \sum_{j \in \mathcal{B}} (1 + \beta) u \left(\frac{Re_j - \tau(e_j) + \psi R \sum_{i=1}^n \tau(e_i) \pi_i}{1 + R} \right) \pi_j \end{aligned}$$

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

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

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

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

Manipulating these equations, and using the assumption that $\beta R = 1$, we get:

$$u'(c_j^y) = \sum_{q=1}^{\bar{k}} \psi u'(c_j^o) \pi_q + \sum_{q=\bar{k}+1}^n \psi u'(c_j) \pi_q, \quad j \notin \mathcal{B}$$

$$u'(c_j) = \sum_{q=1}^{\bar{k}} \psi u'(c_j^o) \pi_q + \sum_{q=\bar{k}+1}^n \psi u'(c_j) \pi_q, \quad j \in \mathcal{B}.$$

Hence, 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. □

3. The optimal tax schedule τ^* equates consumption when young across all agents.

Proof. Using the FOCs described above, we can conclude that $u'(c_j^y) = u'(c_k)$ for $j \notin \mathcal{B}$ and $k \in \mathcal{B}$. This implies, again using the properties of u , that consumption when young is equal across all agents in the economy. □

4. If all agents have positive savings given τ^* , then the optimal tax schedule τ^* is progressive.

Proof. If all agents have positive savings then $\mathcal{B} = \{1, 2, \dots, n\}$ and hence consumption (when young and old) is equal across all agents in the economy. But then, $c_j = c_k$ for $j \neq k$ implies that $e_j - \tau(e_j) = e_k - \tau(e_k)$, which implies that the marginal tax rate for all agents is one, which is a non-decreasing function. Hence, this tax schedule is progressive. □

5. The optimal tax schedule τ^* implies zero informality and zero evasion.

Proof. This is immediate by [Lemma 1](#). □

B Hidden Savings Model Proofs

B.1 Proof of [Lemma 2](#).

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:

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

The structure of the problem arises from composing a strictly concave utility function u with two linear functions of b^c , one (potentially) non-linear function of b , and an additional linear term in b . To ensure that the overall problem is strictly concave, which guarantees a unique solution and the sufficiency of the first-order conditions, it suffices to verify that the composition involving the non-linear function of b preserves concavity. Specifically, it is enough to show that the second derivative of the objective function with respect to b is negative.

Let $F(b, b^c)$ be the objective function. Then:

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

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

Given that $\tau''(b) + (1 + \tau'(b))^2 \geq 0$ and that $-u'(c) < u''(c)$ (these are [Assumption 3](#) and [Assumption 4](#)) then $F_{bb} < u''(c_y) [\tau''(b) + (1 + \tau'(b))^2] \leq 0$, implying that the second derivative of the objective function is indeed negative. \square

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

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

Proof. The first-order conditions that fully characterize agent e 's problem are:

$$u'(c_y)(1 + \tau'(b)) = u'(c_o) + \lambda_1,$$

$$u'(c_y) = \beta u'(c_o) + \lambda_2,$$

$$\lambda_1 \geq 0, \quad b \geq 0, \quad \lambda_1 b = 0,$$

$$\lambda_2 \geq 0, \quad b^c \geq 0, \quad \lambda_2 b^c = 0.$$

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

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

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

Proof. Again, using the FOCs, 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 $\tau'(0) \geq R - 1$. \square

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

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

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

$$u'(c_y)(1 + \tau'(b)) = u'(c_0), u'(c_y) = \beta u'(c_0),$$

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 $\|\cdot\| : \mathcal{T} \rightarrow \mathbb{R}$ 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 5. *The correspondence $\mathcal{G} : \mathcal{T} \rightrightarrows \mathbb{R}_+$ is non-empty and it is actually a function.*

Proof. Let $\tau \in \mathcal{T}, g > 0$. First, I show that $(b_{\tau,g}(e), b_{\tau,g}^c(e))$ are a function of g and that they are both continuous in g . Let us consider the relaxed problem:

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

which ignores the non-negativity constraints imposed to savings. As in the constrained problem, since $\tau \in \mathcal{T}$, the first-order conditions are necessary and sufficient to characterize the solution. Hence, given τ, g , there is only one solution to the problem, which means that $(b_{\tau,g}(e), b_{\tau,g}^c(e))$ are a function of g . The first order-condition with respect to b is given by:

$$F_\tau(b, g) = -u'(c_y)(1 + \tau'(b)) + u'(c_0) = 0.$$

Hence:

$$\begin{aligned} F_b(b, g) &= -u'(c_y)\tau''(b) + u''(c_y)(1 + \tau'(b))^2 + Ru''(c_0), \\ F_g(b, g) &= -\psi u''(c_0). \end{aligned}$$

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

$$F_b(b, g) \leq u''(c_y) \left[\tau''(b) + (1 + \tau'(b))^2 \right] + Ru''(c_0) < 0, \quad (\text{B5})$$

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

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


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

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 $b_{\tau, g}^u(e)$ be the (unconstrained) solution to the agent's problem. The solution to the constrained problem is then:

$$b_{\tau, g}(e) = \max \{b_{\tau}^u(e), 0\},$$

which is also continuous, and differentiable almost for all g except at the \tilde{g} such that $b_{\tau, \tilde{g}}^u(e) = 0$ (which is unique since this is a strictly decreasing function of g). In any case, $b_{\tau, g}(e)$ is a non-increasing continuous function of g . Figure XII plots this construction.

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



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

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

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 $PGC(g)$ is bounded, since τ is a bounded function. Hence $\lim_{g \rightarrow \infty} G_{\tau}(g) = -\infty$. 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 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:

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

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 $\tau'(b_{\tau,g}(e_i)) \frac{b_{\tau,g}(e_i)}{\partial g}$ can be at most one (but not exactly equal to one), which means that:

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

3. **Case 3:** Any other case.

If there are some b 's for which $\tau'(b) = 0$ or $\tau'(b)$ changes from being positive to negative, the same result holds.

□

Lemma 6. 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 $\tau, \tilde{\tau}$ 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 $PGC_\tau(g) = \sum_{i=1}^n \tau(b_{\tau,g}(e_i)) \pi_i$ 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 XIII 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 $G_{\tilde{\tau}}(\mathcal{G}(\tilde{\tau}))$ is 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 XIII presents a graph showing this argument. In conclusion, $\mathcal{G}(\tau)$ is continuous.

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



□

2. The function $b.(\tilde{e}) : \mathcal{T} \rightarrow \mathbb{R}_+$ given by:

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

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 . □

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:

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

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

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

Hence:

$$\begin{aligned} F_b &= -u'(c_y)\tau''(b) + u''(c_y)(1 + \tau'(b))^2 + Ru'(c_o), \\ F_e &= -u''(c_y)(1 + \tau'(b)), \\ F_g &= \psi u''(c_o). \end{aligned}$$

Given [Assumption 3](#) and [Assumption 4](#) then $F_b < 0$, while $F_g > 0$ always. This implies:

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

Now, $F_e > 0$ if $1 + \tau'(b) > -1$. Hence, if this holds:

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

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) $c_{\tau,g}^y(e) > c_{\tau,g}^0(e)$ if and only if $\tau'(b_{\tau,g}(e)) > 0$.

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

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

If $\tau'(b) > 0$ then $u'(c_y) > u'(c_o)$, and by [Assumption 2](#) then $c_y < c_o$. Conversely, since $c_y < c_o$ then $0 > u'(c_y) > u'(c_o)$, implying:

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

yielding $\tau'(b) > 0$.

□

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

Proof. The FOC is the same in this case, leading to the same proof (with some minor changes) as in (a). The same goes for case (c). □

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

2. If $b_{\tau,g}^c(e) > 0$ then $c_{\tau,g}^y(e) > c_{\tau,g}^0(e)$.

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 $u'(c_y) < u'(c_0)$, which by [Assumption 2](#) leads to $c_y > c_0$. □

B.5 Proof of [Proposition 2](#).

Proof. Suppose $\tau \in \mathcal{T}$ is such that $b_{\tau,g}(e) > 0$ and $b_{\tau,g}^c(e) > 0$ for some $e \in E$ and $g = \mathcal{G}(\tau)$. I construct $\hat{\tau} \in \mathcal{T}$ and $\hat{g} = \mathcal{G}(\hat{\tau})$ such that $b_{\hat{\tau},\hat{g}}(e) > b_{\tau,g}(e)$, $b_{\hat{\tau},\hat{g}}^c(e) < b_{\tau,g}^c(e)$, and $\hat{g} > g$. I show that this leads to a strict welfare increase. This construction appeals to the continuity properties of most of the objects involved in the problem. An explicit construction of $\hat{\tau}, \hat{g}$ can be done in the case restricted to affine taxes, which can be seen in ??.

Since $b, b^c > 0$ for type e , then by [Lemma 2](#), it must be the case that $\tau'(b) = R - 1$. The idea of what follows is that the new tax schedule satisfies $\hat{\tau}(b) < \tau(b)$ and $\hat{\tau}'(b) < \tau'(b) = R - 1$, which incentivizes agent e to save strictly more in the bank and less in the couch the couch. This increases e 's welfare since now she is receiving a strictly larger return $R > 1$ for her bank savings. Now, there are some details here that need to be taken care of; the most important one is when we move from τ to $\hat{\tau}$, every agent will potentially change their saving decisions, and not only that, but the public good provision will change as well. Hence, $\hat{\tau}$ needs to be constructed in a way such that every agent has a utility at least equal to what she received when taxes were τ .

Let $\epsilon > 0$, then by [Lemma 6](#) there exists $\delta > 0$ such that $\|\tau - \hat{\tau}\| < \delta$ implies $|\mathcal{G}(\tau) - \mathcal{G}(\hat{\tau})| < \epsilon$, and $|b_{\tau,\mathcal{G}(\tau)}(e) - b_{\hat{\tau},\mathcal{G}(\hat{\tau})}(e)| < \epsilon$. Consider $\hat{\tau}$ within distance δ of τ to be such that $\hat{\tau}(b_{\tau,\mathcal{G}(\tau)}(e)) < \tau(b_{\tau,\mathcal{G}(\tau)}(e))$, $\hat{\tau}'(b_{\tau,\mathcal{G}(\tau)}(e)) < \tau'(b_{\tau,\mathcal{G}(\tau)}(e)) = R - 1$, and $\hat{g} = \mathcal{G}(\hat{\tau}) \leq g = \mathcal{G}(\tau)$. Since bank savings are strictly decreasing in taxes, then $b_{\hat{\tau},\mathcal{G}(\hat{\tau})}(e) > b_{\tau,\mathcal{G}(\tau)}(e)$ and $b_{\hat{\tau},\mathcal{G}(\hat{\tau})}^c(e) > b_{\tau,\mathcal{G}(\tau)}^c(e)$. Notice that, since $R > 1$ and agent e is saving in the bank more under $\hat{\tau}$, while the provision of public good has not fallen when moving from τ to $\hat{\tau}$, then agent e will have a strictly larger utility under the new tax schedule.

In terms of the other agents, since $\hat{\tau}$ is close to τ , then also the savings decisions and therefore each agent's utility, even under the new provision of the public good. Since $\hat{g} \geq g$, then $\hat{\tau}$ yields either the same or a higher utility to all agents not saving at $b_{\tau,\mathcal{G}(\tau)}(e)$. Overall, this represents then a strict increase in utilitarian welfare, implying τ is not optimal. □

B.6 Proof of [Proposition 3](#).

Proof. Fix an optimal tax schedule $\tau^* \in \mathcal{T}$, and let $g^* = G(\tau^*)$ denote the induced public-good provision. For each type $e_i \in E = \{e_1, \dots, e_n\}$, let

$$(b_i^*, b_i^{c*})$$

denote the equilibrium portfolio choice under (τ^*, g^*) .

The first step in this proof is to show a ‘‘robustness’’ property of the agents' savings choices in the following sense: if a tax schedule $\tilde{\tau}$ ‘‘sufficiently’’ approximates τ^* within an interval of each agent's savings choices under τ^* , then those same savings choices should be optimal under $\tilde{\tau}$. This is formalized in the following lemma.

Lemma 7. For each $i \in \{1, \dots, n\}$, there exist numbers $\varepsilon_i > 0$ and $\eta_i > 0$ such that for any tax schedule $\tau \in \mathcal{T}$ satisfying

$$\sup_{b \in (b_i^* - \varepsilon_i, b_i^* + \varepsilon_i)} \left(|\tau(b) - \tau^*(b)| + |\tau'(b) - \tau^{*\prime}(b)| \right) < \eta_i, \quad (\text{B7})$$

the savings portfolio choice of type e_i under (τ, g^*) coincides with $(b_i^*, b_i^{g^*})$.

The proof of Lemma 7 can be found in this appendix, immediately after the proof of this proposition. Now, define the set of equilibrium bank-saving choices:

$$B^* := \{b_i^* : b_i^* > 0\} \subset [0, e_n].$$

Because E is finite, B^* is finite. Let

$$0 < \bar{b}_1 < \bar{b}_2 < \dots < \bar{b}_K$$

denote the *distinct* elements of B^* . For each $k \in \{1, \dots, K\}$, define the set of types choosing that bank-saving level by

$$\mathcal{I}_k := \{i \in \{1, \dots, n\} : b_i^* = \bar{b}_k\}.$$

Now, by Lemma 7, one can choose a single $\varepsilon > 0$ small enough so that: (i) for every distinct positive equilibrium bank-saving level defined above, the closed intervals of radius ε are pairwise disjoint, and (ii) for every type i , the interval around b_i^* of radius ε lies within a robustness neighborhood as above (possibly after taking the minimum over relevant ε_i 's). Define the protected intervals

$$J_k := [\bar{b}_k - \varepsilon, \bar{b}_k + \varepsilon], \quad k = 1, \dots, K,$$

which are pairwise disjoint by construction of ε .

I now proceed to construct $\tilde{\tau} \in \mathcal{T}$ that is concave and that preserves the same savings behavior as τ^* in four substeps.

(i) Pin down $\tilde{\tau}$ on protected intervals. In the case that in each J_k interval τ^* is concave, then simply define:

$$\tilde{\tau}(b) = \tau^*(b) \quad \text{for all } b \in \bigcup_{k=1}^K J_k. \quad (\text{B8})$$

In particular, $\tilde{\tau}$ and τ^* (and their derivatives) coincide on each J_k . In the case that τ^* is not concave in some interval J_k we can construct a concave tax that is "very close" to τ^* but is concave in J_k . What is needed is that the planner can replace τ^* by a concave schedule that is sufficiently close to τ^* in the C^1 sense on small neighborhoods of the finitely many equilibrium choices $\{b_i^*\}$ so that each type's optimal portfolio remains unchanged when $g = g^*$ is held fixed.

Formally, fix $\varepsilon > 0$. Because the set $\{b_i^*\}_{i=1}^n$ is finite, there exist disjoint intervals $\{J_k\}_{k=1}^K$ around the distinct positive equilibrium bank-saving levels and, for each i , robustness constants (ε_i, η_i) from Lemma 7. Choose the J_k 's so that each J_k is contained in the corresponding robustness neighborhood and let $\eta := \min_i \eta_i$. Then one can construct a concave $\hat{\tau} \in T$ such that:

$$\sup_{b \in \bigcup_{k=1}^K J_k} \left(|\hat{\tau}(b) - \tau^*(b)| + |\hat{\tau}'(b) - \tau^{*\prime}(b)| \right) < \eta,$$

and similarly on $[0, \epsilon_i]$ for types with $b_i^* = 0$.

(ii) Linear interpolation on the gaps. List the endpoints of the protected intervals in increasing order:

$$0 = a_0 < a_1 < a_2 < \cdots < a_{2K} < a_{2K+1} = e_n,$$

where for each $k = 1, \dots, K$, $J_k = [a_{2k-1}, a_{2k}]$. Define the gaps

$$G_k := [a_{2k}, a_{2k+1}], \quad k = 0, 1, \dots, K.$$

On each gap G_k , define $\tilde{\tau}$ to be the chord connecting the endpoint values:

$$\tilde{\tau}(b) := \tau^*(a_{2k}) + \frac{\tau^*(a_{2k+1}) - \tau^*(a_{2k})}{a_{2k+1} - a_{2k}} (b - a_{2k}), \quad b \in G_k. \quad (\text{B9})$$

Together with (B8), this defines a continuous, piecewise-linear function on $[0, e_n]$ that matches τ^* on all protected intervals and matches τ^* at every endpoint a_m .

(iii) Enforce concavity by explicit slope pooling. Let $z_0 < z_1 < \cdots < z_M$ be the ordered list of all breakpoints of $\tilde{\tau}$ (these are the a_m 's and possibly additional points if τ^* has breakpoints inside J_k , though under $\tau^* \in C^2$ there are none). On each interval $[z_{m-1}, z_m]$, $\tilde{\tau}$ is linear, so define its slope by

$$s_m := \frac{\tilde{\tau}(z_m) - \tilde{\tau}(z_{m-1})}{z_m - z_{m-1}}, \quad m = 1, \dots, M.$$

A continuous piecewise-linear function is concave if and only if these slopes are weakly decreasing:

$$s_1 \geq s_2 \geq \cdots \geq s_M.$$

If this inequality fails, apply the following pooling procedure (the pool-adjacent-violators algorithm): whenever $s_m < s_{m+1}$ for some m , replace the two adjacent segments $[z_{m-1}, z_m]$ and $[z_m, z_{m+1}]$ by a single segment $[z_{m-1}, z_{m+1}]$ with slope equal to the length-weighted average

$$\bar{s} := \frac{(z_m - z_{m-1})s_m + (z_{m+1} - z_m)s_{m+1}}{z_{m+1} - z_{m-1}}.$$

Iterate this replacement until all slopes are weakly decreasing. The procedure terminates in finitely many steps since each pooling reduces the number of segments by one. Let $\hat{\tau}$ denote the resulting pooled piecewise-linear function. By construction: (a) $\hat{\tau}$ is concave (its slopes are weakly decreasing); (b) $\hat{\tau}(z_m) = \tilde{\tau}(z_m)$ for all breakpoints z_m (pooling preserves endpoint values); (c) in particular, $\hat{\tau}(a_m) = \tilde{\tau}(a_m) = \tau^*(a_m)$ for all m , so $\hat{\tau}$ matches τ^* at all protected-interval boundaries.

(iv) Smooth to ensure $\hat{\tau} \in \mathcal{T}$ while preserving concavity and pinning. The function $\hat{\tau}$ is concave and continuous but may fail to be C^2 at finitely many kink points (the pooled breakpoints). One can remove these kinks by local smoothing *only inside the gaps* $\{G_k\}_{k=0}^K$, leaving $\hat{\tau}$ unchanged on each protected interval J_k . Fix a kink point c of $\hat{\tau}$ with $c \in \text{int}(G_k)$ for some k . Let $s_L \geq s_R$ denote the left and right slopes of $\hat{\tau}$ at c (the inequality holds because $\hat{\tau}$ is concave). Choose $\delta > 0$ small enough that $[c - \delta, c + \delta] \subset \text{int}(G_k)$ and does not intersect any protected interval. Define a quadratic function q on $[c - \delta, c + \delta]$ by imposing the four boundary conditions

$$q(c - \delta) = \hat{\tau}(c - \delta), \quad q'(c - \delta) = s_L, \quad q(c + \delta) = \hat{\tau}(c + \delta), \quad q'(c + \delta) = s_R.$$

There exists a unique quadratic satisfying these four linear constraints. Its second derivative is constant and equal to

$$q'' = \frac{s_R - s_L}{2\delta} \leq 0,$$

so q is concave. Replace $\hat{\tau}$ by q on $[c - \delta, c + \delta]$ and leave it unchanged elsewhere. Perform this operation at each kink point (there are finitely many), choosing the corresponding δ 's small enough so that the smoothing neighborhoods are disjoint and stay inside gaps. The resulting function, still denoted $\hat{\tau}$, is concave, coincides with τ^* on each J_k , and is C^1 on $[0, e_n]$. A second application of the same idea (e.g., using local cubic patches that match second derivatives at endpoints) yields a concave C^2 function. Denote the final smoothed concave schedule by $\tilde{\tau} \in \mathcal{T}$. By construction,

$$\tilde{\tau}(b) = \tau^*(b) \quad \text{for all } b \in \bigcup_{k=1}^K J_k, \quad (\text{B10})$$

and $\tilde{\tau}$ is concave on $[0, e_n]$.

To rule out the possibility that modifying the tax schedule outside the protected intervals generates new globally optimal savings deviations, I complete the construction of $\tilde{\tau}$ so that it never makes any off-equilibrium bank-saving level strictly more attractive than under τ^* . Specifically, after imposing $\tilde{\tau} = \tau^*$ on $\bigcup_{k=1}^K J_k$ (matching both levels and derivatives on those intervals), define $\tilde{\tau}$ on each gap $G_k = [a_{2k}, a_{2k+1}]$ as the smallest concave extension consistent with the pinned endpoint values and satisfying $\tilde{\tau}(b) \geq \tau^*(b)$ for all $b \in G_k$. Such a construction is feasible because concavity is imposed only across gaps while the function is fixed on $\bigcup_k J_k$. By construction, for any type e_i and any bank-saving choice $b \notin \bigcup_k J_k$, the value of the objective under $(\tilde{\tau}, g^*)$ is weakly lower than under (τ^*, g^*) , since the only change operates through the term $-\tilde{\tau}(b)$ and $\tilde{\tau}(b) \geq \tau^*(b)$ at those b . At the candidate equilibrium choices $b_i^* \in \bigcup_k J_k$, the schedules coincide locally, so the objective value is unchanged. Therefore, no type can obtain a strictly higher payoff by deviating to a distant saving level under $\tilde{\tau}$, and by [Lemma 7](#) the portfolio choices (b_i^*, b_i^{c*}) remain globally optimal for each type under $(\tilde{\tau}, g^*)$.

Finally, I show that the equilibrium allocation and public good are unchanged, hence $\tilde{\tau}$ is also optimal. The claim is that the equilibrium portfolio choice of every type under $(\tilde{\tau}, g^*)$ coincides with (b_i^*, b_i^{c*}) . Fix i . If $b_i^* > 0$, then $b_i^* = \bar{b}_k$ for some k , so $b_i^* \in J_k$. By (B10), $\tilde{\tau}$ and τ^* coincide (in particular in C^1) on J_k , which contains a neighborhood of b_i^* . Therefore, the robustness property implies that type i 's portfolio choice is unchanged under $(\tilde{\tau}, g^*)$. If $b_i^* = 0$, then I have shown that there exists $\varepsilon_i > 0$ such that $[0, \varepsilon_i]$ is a robustness neighborhood for the couch choice. By construction of ε and of $\tilde{\tau}$ (which equals τ^* in a neighborhood of 0 whenever 0 is an equilibrium bank-saving boundary, and in any event can be taken C^1 -close near 0 because I did not alter the schedule at 0), one has that $\tilde{\tau}$ is C^1 -close enough to τ^* on $[0, \varepsilon_i]$ to apply the robustness property, so the couch choice is preserved. Thus, for all i ,

$$(b_i(\tilde{\tau}, g^*), b_i^c(\tilde{\tau}, g^*)) = (b_i^*, b_i^{c*}).$$

Since every type's equilibrium portfolio is unchanged, aggregate tax revenue (computed at realized choices) is unchanged. Therefore the government budget constraint implies that the induced public good is the same:

$$G(\tilde{\tau}) = G(\tau^*) = g^*.$$

Welfare depends on τ only through the induced equilibrium allocation and the resulting public

good. Hence welfare under $\tilde{\tau}$ equals welfare under τ^* . Since τ^* is optimal, $\tilde{\tau}$ is also optimal.

In conclusion, I have constructed an optimal tax schedule $\tilde{\tau} \in \mathcal{T}$ that is concave. Hence, without loss of optimality, the planner may restrict attention to concave tax schedules. \square

B.7 Proof of Lemma 7

Proof. Fix i and consider two cases.

Case 1: $b_i^* > 0$. Since $b_i^* > 0$ and there is no evasion (Proposition 2), type e_i saves only in the bank at equilibrium, i.e. $b_i^{c*} = 0$. Lemma 2 implies that mixing between bank and couch occurs only at the knife-edge $\tau'(b) = R - 1$ (at an interior bank choice). Because e_i does *not* mix and chooses an interior bank amount $b_i^* > 0$, we must have

$$\tau^{*'}(b_i^*) \neq R - 1. \quad (\text{B11})$$

Define the strictly positive distance from the knife-edge

$$\delta_i := \frac{1}{2} |\tau^{*'}(b_i^*) - (R - 1)| > 0.$$

By continuity of $\tau^{*'}$ (since $\tau^* \in C^2$), there exists $\varepsilon_i^1 > 0$ such that

$$|b - b_i^*| < \varepsilon_i^1 \implies |\tau^{*'}(b) - \tau^{*'}(b_i^*)| < \delta_i. \quad (\text{B12})$$

Combining (B12) with the definition of δ_i yields, for all b with $|b - b_i^*| < \varepsilon_i^1$,

$$|\tau^{*'}(b) - (R - 1)| \geq |\tau^{*'}(b_i^*) - (R - 1)| - |\tau^{*'}(b) - \tau^{*'}(b_i^*)| > 2\delta_i - \delta_i = \delta_i, \quad (\text{B13})$$

so the knife-edge $\tau'(b) = R - 1$ is uniformly avoided in a neighborhood of b_i^* .

Next, because the individual's objective is strictly concave in the (bank-saving) choice variable on the relevant region, b_i^* is a strict local maximizer of the type- i objective under (τ^*, g^*) . Therefore there exists $\varepsilon_i^2 \in (0, \varepsilon_i^1)$ such that

$$b \in (b_i^* - \varepsilon_i^2, b_i^* + \varepsilon_i^2) \setminus \{b_i^*\} \implies U_i(b; \tau^*, g^*) < U_i(b_i^*; \tau^*, g^*), \quad (\text{B14})$$

where $U_i(\cdot; \tau, g)$ denotes type i 's objective as a function of bank savings given (τ, g) (with the optimal instrument choice embedded; equivalently, one can fix the instrument to be bank in this case since we are away from the knife-edge). Now choose $\varepsilon_i := \varepsilon_i^2$. By continuity of $U_i(\cdot; \tau, g^*)$ in (τ, τ') under the C^1 topology on $[0, e_n]$ (which follows from the smoothness assumptions and the envelope arguments used in the paper), there exists $\eta_i > 0$ such that whenever (B7) holds, the strict inequality (B14) continues to hold with τ in place of τ^* . Consequently, b_i^* remains the unique maximizer of type i 's bank-saving problem in the neighborhood, and in addition, by (B13) and (B7), we also have

$$|\tau'(b) - (R - 1)| > \delta_i/2 > 0 \quad \text{for all } b \in (b_i^* - \varepsilon_i, b_i^* + \varepsilon_i),$$

so the bank-vs-couch comparison remains strict and no mixing can arise (Lemma 2). Hence, type i continues to choose the bank instrument and the same bank savings b_i^* .

Case 2: $b_i^* = 0$. In this case, $b_i^{c*} > 0$ (unless the type saves zero in both instruments, in

which case the argument is trivial). **Lemma 2** implies that $b = 0$ is strictly dominated (at the margin) by couch savings whenever $\tau'(0) > R - 1$. Since type i strictly chooses the couch at equilibrium, we must have

$$\tau^{*'}(0) > R - 1. \quad (\text{B15})$$

Let $\delta_i := \frac{1}{2}(\tau^{*'}(0) - (R - 1)) > 0$. By continuity of $\tau^{*'}$, there exists $\varepsilon_i > 0$ such that

$$b \in [0, \varepsilon_i] \implies \tau^{*'}(b) \geq R - 1 + \delta_i. \quad (\text{B16})$$

If τ satisfies **(B7)** (with this ε_i and some $\eta_i \leq \delta_i/2$), then for all $b \in [0, \varepsilon_i]$,

$$\tau'(b) \geq \tau^{*'}(b) - |\tau'(b) - \tau^{*'}(b)| \geq (R - 1 + \delta_i) - \delta_i/2 = R - 1 + \delta_i/2.$$

Thus, the bank is strictly dominated at the margin throughout $[0, \varepsilon_i]$, and **Lemma 2** implies the type continues to choose the couch instrument (and in particular keeps $b_i = 0$).

In both cases, we have exhibited $\varepsilon_i > 0$ and $\eta_i > 0$ such that **(B7)** guarantees that type i 's equilibrium portfolio choice is unchanged. This completes the proof of the robustness claim. \square

B.8 Proof of Proposition 4.

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

Proof. I first show that if τ^* 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 τ^* . That is, there must exist at least one agent who saves a strictly positive amount $b > 0$ such that $\tau^{*'}(b) = 0$.

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 $b_i^* > 0$, it holds that $\tau^{*'}(b_i^*) > 0$. 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 $b_{\bar{e}}^* = \tilde{b}$.

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

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

Such a construction is feasible within the set \mathcal{T} .

Since type \bar{e} saves $\tilde{b} > 0$ at the optimum and $\tau^{*'}(\tilde{b}) > 0$, 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 $\hat{\tau}$.

For any other type e_i with $b_i^* < \tilde{b}$, the tax schedule $\hat{\tau}$ coincides with τ^* in a neighborhood of b_i^* . By **Lemma 7**, their equilibrium portfolio choices remain unchanged. Moreover, by the envelope theorem, their utilities vary only at second order. Hence, for ε sufficiently

small, their welfare does not decrease.

Therefore, aggregate welfare under $\hat{\tau}$ is strictly higher than under τ^* , contradicting the optimality of τ^* . 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 $\tau^{*'}(\bar{b}) = 0$.

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

2. At τ^* there is zero evasion.

Proof. This is a direct consequence of [Proposition 2](#). \square

3. 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 $b > \tilde{b}$. \square

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 $X \in \{0, 1, \dots, n\}$ 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 XIV](#).

Figure XIV: Beta-Binomial Density Function.



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 $(b_{\tau,g}(e), b_{\tau,g}^c(e))$.

Using these optimal savings decisions, I then compute the implied public good contribution as a function of τ and g , which I denote by $\text{PGC}(\tau, g) = R \sum_{e \in E} \tau(b_{\tau,g}(e)) \pi(e)$. 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 standard root-finding techniques applied to the function $\text{PGC}(\tau, g) - g$.

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 . Second, τ must lie in the set \mathcal{T} , which imposes bounds on both the first and second derivatives of the tax schedule. These constraints are handled directly within the optimization algorithm to ensure that the resulting function is both implementable and consistent with the assumptions laid out in [Section 5](#).

Finally, I exploit the structure of the optimal tax schedule derived in [Section 5](#). In particular, I consider an additional step in which the tax function τ is defined as a hybrid object: for $b > \tilde{b}$, $\tau(b) = \bar{\tau}$ is constant, while for $b \leq \tilde{b}$, $\tau(b)$ is approximated by a Chebyshev polynomial. To ensure smoothness at the junction point $b = \tilde{b}$, I impose that the polynomial satisfies $\tau(\tilde{b}) = \bar{\tau}$ and $\tau'(\tilde{b}) = 0$. This formulation both respects the theoretical structure of the optimal tax and improves numerical precision by reducing the dimensionality of the function space over which optimization is performed.