Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Access Frictions, Optimal Redistribution, and the Bias in Standard Tax Rate Computations

Anirban Ghatak

aghatak@iimk.ac.in

Indian Institute of Management Kozhikode

Research Article

Keywords: Optimal income taxation, access frictions, sufficient statistics, administrative burdens, participation, joint optimum, calibration

Posted Date: March 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9142849/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Access Frictions, Optimal Redistribution, and the Bias in Standard Tax Rate Computations

Anirban Ghatak*

Abstract

Standard optimal income tax theory assumes that disposable resources translate directly into individual welfare. In reality, accessing markets, public services, and transfer programs requires effort, information, and resources that are themselves scarce and more scarce at the bottom of the income distribution. We introduce *access frictions* into a Mirrleesian framework by modeling a smooth wedge between disposable income and effective consumption. The wedge is governed by an intensity index $\pi(a, c)$ that depends on the household's disposable income c and on a costly government instrument a that captures the quality of service delivery, administrative design, and program accessibility. We characterize optimal income taxation in this environment through four sets of results. First, a Saez-style top-bracket formula in which frictions enter through two channels: welfare weights are rescaled by the marginal effectiveness $m(a, c) = \partial\tilde{c}/\partial c$, and behavioral responses are governed by the elasticity of earnings with respect to the composite incentive shifter $s = (1 - T') \cdot m$ rather than the statutory net-of-tax rate alone. Second, under quasi-linear preferences, a full-distribution formula that nests the top-bracket result and shows that benchmark computations abstracting from frictions overstate optimal marginal rates at every income level. Third, a participation tax formula in the spirit of Saez [2002], showing that access frictions attenuate the effective participation gain and bias standard Earned Income Tax Credit (EITC)-style analyses in a direction that depends on whether frictions are more severe in employment or non-employment. Fourth, a characterization of the jointly optimal tax schedule and access policy, establishing that the optimal level of access investment rises with income inequality under a mild curvature condition consistent with standard parameterizations of preferences. We complement the theoretical analysis with a calibration exercise that pins down the friction gradient $\gamma = m_c \cdot c/m$ from experimental and survey evidence on administrative burdens and program take-up. Across a wide range of standard parameter values, the friction correction reduces the implied optimal top marginal tax rate by 1.5 to 4.5 percentage points, and compresses the U-shaped rate schedule most strongly at the bottom of the income distribution, where frictions are steepest.

Keywords— Optimal income taxation, access frictions, sufficient statistics, administrative burdens, participation, joint optimum, calibration

*Indian Institute of Management Kozhikode, Kerala, India, 673570. Email: `aghatak@iimk.ac.in`. I use the plural throughout purely as a matter of expository convention.

1 Introduction

Standard optimal income tax theory inhabits a frictionless world in which a dollar of after-tax income is a dollar of consumption. From Mirrlees [1971] through the sufficient-statistics reformulations of Diamond [1998] and Saez [2001], disposable resources translate directly into individual welfare, and the policy problem reduces to choosing a statutory schedule that balances redistributive gains against the distortionary cost of discouraging work. This frictionless benchmark is analytically powerful. It abstracts from a feature of redistribution systems that is both empirically pervasive and theoretically consequential: households must spend resources (time, information, psychological bandwidth, legal status, social connections) to convert disposable income into effective consumption.

The empirical record on this point is substantial. Finkelstein and Notowidigdo [2019] find that information and application costs deter roughly forty per cent of eligible non-participants from claiming Supplemental Nutrition Assistance Program (SNAP) benefits, even when the financial gain from doing so is large. Ko and Moffitt [2022] survey a large literature documenting take-up rates well below one for programs ranging from Medicaid to housing assistance to earned income tax credits. Herd and Moynihan [2018] argue that compliance costs, documentation requirements, and psychological costs function as de facto policy instruments that shape who receives what, often in ways that are regressive: precisely the households most in need of transfers bear the highest costs of accessing them. Importantly, these frictions are not fixed. They depend on household resources: richer, more educated households navigate complex systems at lower cost. They also respond to deliberate choices about administrative design, program complexity, and service delivery quality. A government that simplifies its application process or expands outreach is reducing access frictions; one that increases documentation requirements or reduces caseworker staffing is raising them.

This paper introduces *access frictions* into the Mirrleesian optimal tax framework as a structural feature of the welfare-production technology. The key modeling element is an *intensity index* $\pi(a, c) \in (0, 1)$ that governs the share of the potential consumption gain above an outside option that a household can actually realize. Effective consumption is

$$\tilde{c} = \underline{c} + \pi(a, c)\,(c - \underline{c}),$$

where $\underline{c}$ is an outside-option consumption floor, $c = y - T(y)$ is disposable income, and $a \geq 0$ is a government instrument that shifts access quality at a resource cost κa. The function $\pi(a, c)$ is increasing in a (better administrative design improves access) and non-decreasing in c (higher-income households realize a weakly larger fraction of their potential consumption). The resulting object $m(a, c) \equiv \partial\tilde{c}/\partial c$ is the *marginal effectiveness* of disposable income, measuring how many cents of effective consumption an additional cent of disposable income generates.

The central insight is that $m(a, c)$ enters optimal tax formulas through two distinct channels. First, the social value of redistributing a dollar from high earners to low earners is not just $u'(c)$ but $u'(\tilde{c}) \cdot m(a, c)$. Even if two households have the same disposable income, the household facing heavier frictions gains less effective consumption from a marginal transfer, which reduces the welfare weight the planner assigns to it. Second, the behavioral return to earning an additional dollar is not the statutory net-of-tax rate $n = 1 - T'$ but the composite shifter $s = n \cdot m$. A high-earner who faces low frictions has a stronger effective incentive to work than a low-earner facing high frictions, even if both face the same marginal tax rate. Both channels modify the classical tradeoff between redistribution and efficiency in ways that depend systematically on the joint distribution of income and access quality across the population.

We develop these implications through four sets of theoretical results and one calibration exercise. The first is a Saez-style top-bracket characterization (Theorem 1): the optimal top marginal tax rate satisfies

the standard inverse-elasticity formula, with welfare weights inflated by $m(a, c)$ and the behavioral elasticity defined with respect to $s = n \cdot m$ rather than n alone. Under the empirically plausible condition that m is increasing in c (so that effective incentives co-move with income), the total elasticity e_s exceeds the partial elasticity e_n, and benchmark computations that use e_n in place of e_s systematically overstate the optimal top rate (Theorem 3). The second set of results extends this characterization to the entire income distribution under quasi-linear preferences (Theorem 2), showing that the over-taxation bias extends to every income level and that the friction-adjusted rate schedule is uniformly below the frictionless benchmark when $e_s > e_n$ throughout (Corollary 1). The third result concerns labor-market participation: when households face a binary decision whether to enter the labor market, access frictions modify the optimal participation tax by rescaling both the social welfare weight and the participation elasticity through the friction pass-through m (Proposition 2). This generates a directional bias in standard EITC-style analyses whose sign depends on whether frictions are heavier for employed or non-employed households, a distinction that varies systematically across program types and institutional contexts. The fourth result characterizes the jointly optimal tax schedule and access policy (Theorem 4) and establishes that the optimal investment in access improvement rises with income inequality under a curvature condition on utility that is satisfied under standard constant relative risk aversion (CRRA) parameterizations with relative risk aversion above one (Proposition 5).

The calibration exercise (Section 12) translates the theoretical bias result into quantitative magnitudes. We pin down the friction gradient $\gamma \equiv m_c \cdot c/m$ from three empirical sources: the experimental evidence of Finkelstein and Notowidigdo [2019] on SNAP application costs, the income gradient of administrative burden documented by Herd and Moynihan [2018], and cross-program take-up evidence in Ko and Moffitt [2022]. Paired with standard estimates of the earnings elasticity, Pareto tail parameter, and welfare weights from Diamond and Saez [2011] and Saez et al. [2012], the calibration yields friction corrections of 1 to 4.5 percentage points on the optimal top rate across all parameter combinations. At the bottom of the income distribution, where frictions are steepest, the correction factor reduces the optimal rate to roughly 77 per cent of the frictionless benchmark. These magnitudes are comparable in scale to the differences between major tax reform proposals, which suggests the friction channel is not merely a theoretical refinement.

The paper's contribution is to bridge the sufficient-statistics tradition in optimal income taxation and the empirical literature on take-up and administrative burdens. The existing sufficient-statistics literature extends optimal tax formulas by expanding the set of behavioral margins [Saez, 2002, Piketty et al., 2014] or introducing screening instruments alongside the tax schedule [Kleven and Kopczuk, 2011]. The take-up literature documents that frictions suppress effective redistribution. That literature has not formally integrated this finding into the optimal tax problem. The present paper does so by placing frictions inside the welfare-production technology rather than treating them as an add-on cost, which allows the sufficient-statistics apparatus to carry over while generating modified formulas with clear empirical content.

The paper is organized as follows. Section 2 discusses related literature in detail. Section 3 introduces the model. Sections 4 and 5 describe the individual and government problems respectively. Section 6 derives the top-bracket formula; Section 7 extends it to the full income distribution. Section 8 establishes the attenuation result and Section 9 the over-taxation bias. Section 10 introduces the participation margin, and Section 11 characterizes the joint optimum and the comparative static with respect to inequality. Section 12 presents the calibration. Section 13 discusses policy implications and limitations. Section 14 concludes.

2 Related Literature

This paper contributes to three bodies of literature: optimal income taxation, the economics of administrative burdens and take-up, and the design of targeting instruments in redistribution programs.

Optimal income taxation. The canonical Mirrlees [1971] framework posits a planner who observes earnings but not ability, and designs a nonlinear income tax to balance redistribution against incentive costs. The key empirical advance came with Saez [2001], who showed that the optimal marginal tax rate at any income level can be expressed in terms of three sufficient statistics (the earnings elasticity, the social welfare weight, and the local thickness of the income distribution) without requiring knowledge of the structural preference parameters underlying those objects. Chetty [2009] articulates the general method. Diamond [1998] works out the full-distribution formula under quasi-linear preferences and establishes the U-shaped pattern of optimal marginal rates: high at the bottom (to phase out transfers), low in the middle (where the earnings density is thick), and converging to a Pareto tail formula at the top. The present paper derives analogs of all three of these results (the top-bracket formula, the full-distribution formula, and the U-shape analysis) modified by the presence of access frictions.

Several papers extend the sufficient-statistics approach by introducing additional behavioral margins. Saez [2002] shows that when participation decisions respond to tax incentives on the extensive margin, the optimal policy involves subsidizing employment for low earners, rationalizing the EITC. Piketty et al. [2014] decompose the top-income elasticity into real supply, avoidance, and bargaining components, each of which enters the optimal top-rate formula differently. The present paper is closest in method to this literature. We add a new behavioral margin (the friction pass-through $m(a, c)$) that modifies both the welfare weights and the effective elasticity in the optimal tax formula.

Administrative burdens and take-up. A large empirical literature documents that eligible households frequently fail to claim benefits to which they are entitled. Moffitt [1983] introduced the idea that welfare participation carries a non-pecuniary cost, modeled as a utility loss from stigma. Currie [2006] surveys the evidence and argues that transaction costs (the time, effort, and navigational resources required to apply) are at least as important as stigma, and are more amenable to policy intervention. Finkelstein and Notowidigdo [2019] provide experimental confirmation from SNAP, finding that information-based interventions raise take-up substantially even among households already aware of the program, which suggests that the barrier is not primarily informational. Cross-country patterns are documented by Hernanz et al. [2004], and a comprehensive recent survey is Ko and Moffitt [2022]. Herd and Moynihan [2018] develop the concept of *administrative burden* as a policy instrument in its own right: governments that design programs with high compliance costs are, intentionally or not, restricting access to redistribution.

These empirical findings motivate the present model. The paper differs from the take-up literature in an important respect. Take-up models typically treat access costs as fixed utility losses or fixed resource costs, a "stigma" or "hassle" that is either incurred or not. The present paper models access as a *smooth intensity* that governs the realizable fraction of the consumption gain from receiving income, which allows it to enter the sufficient-statistics calculus in a tractable way.

The smooth formulation is not merely a technical convenience but a substantive modeling choice supported by the structure of the empirical evidence. The fixed-cost tradition, originating with Moffitt [1983] and formalized in Kleven and Kopczuk [2011], treats access costs as lump sums that are either incurred or avoided, generating a participation margin at which take-up jumps discretely. That structure is appropriate when the dominant friction is a stigma cost attached to program receipt as such, or a fixed complexity

cost per application that is independent of household resources. The evidence, however, points toward a different pattern. Finkelstein and Notowidigdo [2019] find that the marginal enrollees induced into SNAP by application assistance have strictly higher net incomes than the average existing enrollee. This income composition result is inconsistent with a uniform fixed cost, which would not predict any income gradient at the margin; it is consistent with a friction cost that is decreasing and smooth in income, so that higher-income households are disproportionately represented among those who are just barely deterred by the access barrier. Herd and Moynihan [2018] decompose administrative burden into learning costs, compliance costs, and psychological costs, all three of which scale continuously with socioeconomic resources (education, time availability, cognitive bandwidth, and social networks) rather than being incurred as a fixed amount. Currie [2006] documents that reductions in recertification frequency raise SNAP and Medicaid take-up progressively across the income distribution, not as a threshold effect, which is again consistent with a smooth friction function. The present paper therefore models access costs through a continuously differentiable intensity $\pi(a, c)$ that is increasing in disposable income, rather than as a fixed utility loss. This choice also has a direct analytical payoff that is the smooth formulation preserves the differentiability required for the perturbation arguments underlying Theorems 1 and 2, whereas a binary access cost would require a separate treatment of the participation margin along the lines of Saez [2002], which is taken up in Section 10.

Screening and joint instruments. A distinct literature examines how non-tax instruments can be used alongside the income tax to improve targeting. Nichols and Zeckhauser [1982] show that "ordeals" (requirements that impose costs on applicants) can screen out ineligible mimickers when the ordeal cost is higher for high-ability types. Besley and Coate [1992] formalize the analogous argument for workfare: requiring program recipients to perform unproductive labor can deter high-ability individuals from claiming transfers. Both arguments require the ordeal cost to be positively correlated with outside options, a condition that fails precisely when frictions are regressive, as the empirical evidence suggests. Kleven and Kopczuk [2011] provide the most direct antecedent within the optimal tax literature, showing that optimal program complexity involves an interior trade-off between deterring ineligible applicants and imposing costs on the truly needy. In the present paper, access policy a plays a different role. It reduces frictions for all households rather than screening among them, and its optimal level is determined by equating the aggregate social benefit of improved access to its resource cost. The joint optimum section (Section 11) thus complements Kleven and Kopczuk by considering the case in which the friction function is *decreasing* in income (so that frictions hurt the poor more) in place of the Nichols and Zeckhauser case in which they screen out the rich.

Finally, Werquin [2016] introduces a different kind of friction into optimal taxation: fixed adjustment costs in labor supply that generate lumpy earnings responses. Our frictions operate on the consumption-realization side rather than the earnings-generation side, and they generate smooth rather than lumpy responses.

3 The Model

3.1 Population and Earnings Technology

The economy is populated by a continuum of agents, each characterized by a productivity type $\theta \in \Theta \subset \mathbb{R}_+$, drawn from a population with density $f(\theta)$ and cumulative distribution function (CDF) $F(\theta)$. An agent of type θ who supplies effort $e \geq 0$ earns income $y = \theta e$. The type θ is private information; only earnings y are observed by the government. This is the standard Mirrleesian informational structure [Mirrlees, 1971], in

which the planner can condition policy only on observable earnings and not on the underlying ability that generated them.

Given a tax schedule $T : \mathbb{R}_+ \to \mathbb{R}$, disposable income is

$$c \equiv y - T(y).$$

We write $n(y) \equiv 1 - T'(y)$ for the net-of-tax rate at income y, the fraction of a marginal dollar of earnings that the agent retains after tax.

3.2 Access Technology

The key departure from the standard model is that disposable income c does not translate one-for-one into effective consumption. In practice, a household must invest time, effort, and ancillary resources to access markets, claim program benefits, navigate institutions, and convert statutory entitlements into realized welfare. We model this wedge through a smooth intensity function.

Let $\underline{c} \geq 0$ denote an outside-option or floor level of effective consumption available to all agents regardless of their disposable income, for instance the effective consumption achievable through subsistence, informal networks, or unconditional transfers. Effective consumption is defined as the weighted average between $\underline{c}$ and c, where the weight on c is governed by the access intensity:

$$\tilde{c} \;\equiv\; \underline{c} + \pi(a, c)\,(c - \underline{c}). \tag{1}$$

The function $\pi(a, c) \in (0, 1)$ is the *access intensity*, meaning the fraction of the potential consumption gain $(c - \underline{c})$ that the household can actually realize given access quality a and disposable income c. When $\pi = 1$, there are no frictions and $\tilde{c} = c$. When π is strictly below one, some fraction of disposable income is dissipated in the process of converting it into welfare.

The instrument $a \geq 0$ represents government choices that affect access quality, including the staffing and design of administrative offices, the complexity of application procedures, the quality of outreach and information provision, the availability of digital infrastructure, and similar determinants of how easily households can realize the consumption value of their income. A higher a means lower frictions. Providing a costs the government κa in revenue-equivalent units, where $\kappa > 0$ is the marginal resource cost of access improvement.

We focus throughout on the economically relevant region $c \geq \underline{c}$.

Assumption 1 (Access Technology)**.** *The function $\pi(a, c)$ is twice continuously differentiable on $\{(a, c) : a \geq 0,\ c \geq \underline{c}\}$ and satisfies, at every such (a, c):*

$$\pi(a, c) \in (0, 1), \qquad \pi_a(a, c) > 0, \qquad \pi_c(a, c) \geq 0.$$

The sign restrictions in Assumption 1 have direct economic interpretations. The condition $\pi_a > 0$ says that better access policy increases the realized intensity of consumption: administrative improvements actually help. The condition $\pi_c \geq 0$ says that households with higher disposable income realize weakly more of their consumption potential. This captures the empirically documented pattern that access frictions are regressive: higher-income households are better educated, have more time and social capital, and face lower marginal costs of navigating administrative requirements [Herd and Moynihan, 2018, Currie, 2006]. The strict inequality $\pi < 1$ at every point captures the assumption that frictions are always present to some degree, a normalization that can be relaxed without affecting the results qualitatively.

A central object in the analysis is the *marginal effectiveness* of disposable income, defined as the derivative of effective consumption with respect to disposable income:

$$m(a,c) \equiv \frac{\partial \tilde{c}}{\partial c} = \pi(a,c) + \pi_c(a,c)\,(c - \underline{c}). \tag{2}$$

The object $m(a,c)$ measures how many additional units of effective consumption a household gains from one additional unit of disposable income. When $\pi_c = 0$, we have $m = \pi < 1$. An additional dollar of disposable income translates into only π cents of effective consumption because the household can only access a fraction π of what it earns. When $\pi_c > 0$, there is an additional intensive-margin effect. Higher income raises the intensity π itself, so the realized fraction of the marginal dollar is larger than π alone.

Remark 1 (Upper bound on marginal effectiveness)**.** *Assumption 1 guarantees $m(a,c) \geq \pi(a,c) > 0$ but does not impose an upper bound on m. If one interprets m as an attenuation factor (capturing the idea that frictions strictly erode the value of income), one would impose $m(a,c) \leq 1$ as an additional restriction. None of the results below require this bound, and we refrain from imposing it so as to maintain generality. The economically relevant case for our purposes is $m < 1$, and the algebra applies for any value of m.*

Microfoundation for the Access Technology

Consider a household whose intended consumption bundle can be separated into two components. A fraction $1-\lambda$ of the bundle consists of goods and services that the household can access without friction. The household simply spends its income and receives full value. The remaining fraction λ consists of goods or services (publicly provided healthcare, subsidized housing, program benefits, formal financial products) whose effective realization requires navigation of an institutional or administrative process. Conditional on access level a and disposable income c, the household can realize a fraction $\pi(a,c)$ of this access-dependent component, giving

$$\tilde{c} = (1-\lambda)\,c + \lambda[\,\underline{c} + \pi(a,c)\,(c - \underline{c})]\,.$$

Normalizing $\lambda = 1$ (so that all consumption is access-dependent) or equivalently redefining $\underline{c}$ to absorb the frictionless component recovers the specification in equation (1). The parameter $\pi(a,c)$ thus has a natural interpretation as the realized quality or intensity of access-dependent consumption, and Assumption 1 maps onto clear properties of the underlying institutional process.

4 Individual Optimization

Given the tax schedule $T(\cdot)$ and access level a, an agent of type θ chooses effort e (equivalently, earnings $y = \theta e$) to maximize utility. Preferences over effective consumption and effort are represented by $u(\tilde{c}) - v(y/\theta)$, where u is strictly increasing and strictly concave, v is strictly increasing and strictly convex, and both are twice continuously differentiable. The agent takes $T(\cdot)$ and a as given. Using the notation $c(y) \equiv y - T(y)$ and $\tilde{c}(y) \equiv \underline{c} + \pi(a, c(y))(c(y) - \underline{c})$, the agent solves:

$$\max_{y \geq 0}\; u\big(\tilde{c}(y)\big) - v\big(\tfrac{y}{\theta}\big)\,. \tag{3}$$

Lemma 1 (Agent First-Order Condition)**.** *Under Assumption 1, an interior optimum $y(\theta)$ to problem (3)*

satisfies

$$u'\big(\tilde{c}(y)\big) \cdot (1 - T'(y)) \cdot m(a, c(y)) \;=\; \frac{1}{\theta}\, v'\Big(\frac{y}{\theta}\Big), \tag{4}$$

where $m(a, c(y))$ *is as defined in* (2).

Proof. Define $\Phi(y) \equiv u(\tilde{c}(y)) - v(y/\theta)$. The chain rule gives $\Phi'(y) = u'(\tilde{c}(y))\,\tilde{c}'(y) - v'(y/\theta)/\theta$. We compute $\tilde{c}'(y)$ by differentiating $\tilde{c}(y) = \underline{c} + \pi(a, c(y))(c(y) - \underline{c})$:

$$\tilde{c}'(y) = \pi(a, c(y))\, c'(y) + \pi_c(a, c(y))\, c'(y)\, (c(y) - \underline{c}) = c'(y)\big[\pi(a, c(y)) + \pi_c(a, c(y))(c(y) - \underline{c})\big].$$

Since $c'(y) = 1 - T'(y)$ and the bracket equals $m(a, c(y))$ by definition (2), we have $\tilde{c}'(y) = (1 - T'(y))\, m(a, c(y))$. Setting $\Phi'(y) = 0$ yields (4). □

Lemma 1 has a natural economic reading. In the frictionless benchmark, the agent's optimality condition equates the marginal benefit of an additional dollar of earnings ($u'(c)(1 - T')$, the marginal utility of consumption multiplied by the net-of-tax fraction retained) to its marginal cost, $v'(y/\theta)/\theta$. With access frictions, this condition is modified. The effective marginal benefit of earning is $u'(\tilde{c})(1 - T') \cdot m$, because an additional dollar of disposable income converts into only m dollars of effective consumption rather than one. The composite object

$$s(y) \equiv (1 - T'(y)) \cdot m(a, c(y)) \tag{5}$$

is the *effective net-of-tax rate*, the effective consumption return to a marginal dollar of gross earnings, accounting for both the statutory tax wedge and the access friction. This quantity plays the role of the incentive shifter in all of the sufficient-statistics characterizations that follow.

5 The Government Problem

The government operates as a Mirrleesian planner. It observes earnings but not types, and designs the policy pair $(T(\cdot), a)$ to maximize social welfare subject to incentive compatibility and a budget constraint. Social welfare is utilitarian:

$$W(T, a) = \int_{\Theta} \big[u(\tilde{c}(\theta)) - v(e(\theta))\big]\, f(\theta)\, d\theta, \tag{6}$$

where we write $\tilde{c}(\theta) \equiv \tilde{c}(y(\theta), a)$ and $e(\theta) = y(\theta)/\theta$ as shorthands for the effective consumption and effort of type θ at the equilibrium allocation.

Incentive compatibility requires that each type θ prefers the bundle $(y(\theta), T(y(\theta)))$ intended for it over any other bundle available on the tax schedule. The standard approach, which we adopt, is to work with the perturbation arguments of Saez [2001] and Chetty [2009], which characterize necessary conditions for optimality without fully solving for the optimal schedule.

The government's resource constraint must account for both tax revenue and the direct cost of access policy. Net revenue must be non-negative:

$$\int_{\Theta} T(y(\theta))\, f(\theta)\, d\theta \;\geq\; \kappa a. \tag{7}$$

We denote by $\lambda > 0$ the Lagrange multiplier on this constraint, interpreted as the marginal social value of public funds. At any interior solution, λ is determined endogenously by the budget constraint holding with equality.

6 The Top-Bracket Formula

We begin by characterizing the optimal marginal tax rate in the upper tail of the income distribution. The method is the top-bracket perturbation argument of Saez [2001], adapted to the present environment.

6.1 Perturbation and Sufficient Statistics

Fix a threshold income $\bar{y} > 0$. Consider a small reform to the tax schedule that raises the marginal rate by $d\tau > 0$ on all incomes above $\bar{y}$ while leaving the schedule unchanged below. Such a reform affects agents above the threshold in two ways. Mechanically, it collects additional revenue from those agents. Behaviorally, it induces them to earn less in response to the reduced effective return. We assume no bunching at $\bar{y}$ under the baseline schedule, so that all perturbation effects are differentiable in $d\tau$.

Let $h(\cdot)$ and $H(\cdot)$ denote the density and CDF of equilibrium earnings under the baseline policy (T, a). Define the *income-weighted mass above the threshold*:

$$A(\bar{y}) \equiv \int_{\bar{y}}^{\infty} (z - \bar{y})\, h(z)\, dz,$$

which equals the total mechanical tax increase per unit of $d\tau$ on the population above $\bar{y}$.

The key sufficient statistics are the *social marginal welfare weight* at earnings level z,

$$g(z) \equiv \frac{u'(\tilde{c}(z)) \cdot m(a, c(z))}{\lambda}, \tag{8}$$

and the *composite earnings elasticity*

$$e_s(\bar{y}) \equiv \left. \frac{d \log z}{d \log s} \right|_{z=\bar{y}}, \tag{9}$$

where $s(z) = (1 - T'(z)) \cdot m(a, c(z))$ is the effective net-of-tax rate from (5). The welfare weight $g(z)$ measures, in units of public funds, the social value of a marginal dollar of effective consumption delivered to an agent at income z. Crucially, $g(z)$ incorporates $m(a, c(z))$ because frictions reduce the effective value of disposable income, so the welfare weight is scaled down by the marginal effectiveness of transfers. The elasticity e_s captures the total earnings response to a change in the effective incentive to earn, allowing m to move endogenously as disposable income changes.

For comparison, we also define the partial elasticity holding the access intensity fixed:

$$e_n(\bar{y}) \equiv \left. \frac{\partial \log z}{\partial \log n} \right|_{m,\, z=\bar{y}}. \tag{10}$$

The distinction between e_s and e_n will be central to the bias result in Section 9.

6.2 Optimal Top Marginal Rate

Define the *local Pareto parameter*

$$\alpha(\bar{y}) \equiv \frac{\bar{y}\, h(\bar{y})}{1 - H(\bar{y})},$$

the elasticity of the upper-tail mass $1 - H(z)$ with respect to z, evaluated at $\bar{y}$. Together with the income-weighted average welfare weight above the threshold,

$$\bar{g}(\bar{y}) \equiv \frac{1}{A(\bar{y})} \int_{\bar{y}}^{\infty} g(z)\,(z - \bar{y})\,h(z)\,dz,$$

these objects characterize the optimal top rate.

Theorem 1 (Optimal Top Marginal Tax Rate)**.** *Suppose there is no bunching at $\bar{y}$ and the perturbation $d\tau$ is admissible. At any interior optimum, the optimal top marginal tax rate $\tau(\bar{y})$ satisfies*

$$\frac{\tau(\bar{y})}{1 - \tau(\bar{y})} \;=\; \frac{1 - \bar{g}(\bar{y})}{\alpha(\bar{y})\,e_s(\bar{y})}. \tag{11}$$

Proof. We account for three effects of the perturbation, namely the mechanical revenue gain, the mechanical welfare cost, and the behavioral revenue loss.

Mechanical revenue gain. Holding all agents' behavior fixed, the reform raises each agent above $\bar{y}$ with earnings z by an additional $d\tau\,(z - \bar{y})$ in tax liability. Integrating over the earnings distribution:

$$dR_{\text{mech}} = d\tau \int_{\bar{y}}^{\infty} (z - \bar{y})\,h(z)\,dz = d\tau\,A(\bar{y}).$$

Mechanical welfare cost. The reform reduces disposable income of each agent above $\bar{y}$ by $d\tau\,(z - \bar{y})$. Since $\partial\tilde{c}/\partial c = m(a, c(z))$ by definition (2), the reduction in effective consumption for an agent at earnings z is $m(a, c(z)) \cdot d\tau\,(z - \bar{y})$. The social welfare loss, measured in welfare units, is therefore:

$$dW_{\text{mech}} = -\int_{\bar{y}}^{\infty} u'(\tilde{c}(z)) \cdot m(a, c(z)) \cdot d\tau\,(z - \bar{y})\,h(z)\,dz.$$

Dividing and multiplying by λ and using the definition of $g(z)$ in (8):

$$dW_{\text{mech}} = -\lambda\,d\tau \int_{\bar{y}}^{\infty} g(z)\,(z - \bar{y})\,h(z)\,dz = -\lambda\,d\tau\,A(\bar{y})\,\bar{g}(\bar{y}).$$

Behavioral revenue loss. The reform reduces the net-of-tax rate n by $d\tau$ for all agents above $\bar{y}$. This shifts the effective net-of-tax rate $s = n \cdot m$ by $ds = m(a, c(\bar{y})) \cdot (-d\tau)$, so $d\log s \approx -d\tau/n(\bar{y})$ to first order.[1] The earnings response of each agent near $\bar{y}$ is, by definition of e_s:

$$\frac{dz}{z}\bigg|_{z=\bar{y}} = e_s(\bar{y}) \cdot \frac{ds}{s}\bigg|_{z=\bar{y}} = -e_s(\bar{y}) \cdot \frac{d\tau}{n(\bar{y})}.$$

Each agent who reduces earnings by dz causes a revenue loss of $T'(\bar{y})\,dz = \tau(\bar{y})/(1 - \tau(\bar{y})) \cdot (1 - \tau(\bar{y}))\,dz$. Aggregating over the density $h(\bar{y})$ at the threshold and using the definition of $\alpha(\bar{y})$:

$$dR_{\text{beh}} = -d\tau \cdot \tau(\bar{y}) \cdot e_s(\bar{y}) \cdot \alpha(\bar{y}) \cdot A(\bar{y}).$$

[1] This uses $d\log n = -d\tau/n(\bar{y})$ and $d\log m \approx 0$ to first order for a perturbation around $\bar{y}$, since the access intensity adjusts only as earnings respond.

Optimality. At an interior optimum, the perturbation must have zero first-order effect on the Lagrangian:

$$dW_{\text{mech}} + \lambda\big(dR_{\text{mech}} + dR_{\text{beh}}\big) = 0.$$

Substituting and dividing by $\lambda\, d\tau\, A(\bar{y}) > 0$:

$$-\bar{g}(\bar{y}) + 1 - \tau(\bar{y})\, e_s(\bar{y})\, \alpha(\bar{y}) = 0.$$

Solving for $\tau(\bar{y})$ and expressing as a ratio yields (11). □

Remark 2 (Relation to the standard formula)**.** *Theorem 1 is structurally identical to the top-bracket formula in Saez [2001], with two modifications. First, the social marginal welfare weight $g(z)$ in (8) incorporates $m(a, c(z))$, rescaling the utilitarian weight $u'(\tilde{c})/\lambda$ by the marginal effectiveness of disposable income at income z. Second, the behavioral elasticity is e_s, defined with respect to the composite incentive shifter $s = n \cdot m$, rather than the elasticity with respect to n alone. In the frictionless benchmark, $\pi \equiv 1$, yielding $m \equiv 1$, $\tilde{c} = c$, $e_s = e_n$, and (11) reduces exactly to the Saez formula.*

7 The Full-Distribution Formula

Theorem 1 pins down optimal marginal rates only in the upper tail of the earnings distribution. The top-bracket perturbation works by varying the rate above a threshold and examining agents above it, so its conclusions are local to the tail. To characterize optimal rates across the full distribution, we apply an interior perturbation argument: raise the marginal rate on a thin interval $[z^*, z^* + \varepsilon]$ and count both the costs (borne by agents above z^* through reduced incentives) and the benefits (the revenue gain from all agents above z^* who face a higher rate). This interior perturbation generates income effects through the earnings-disutility function v, which complicate the welfare analysis. We eliminate these effects by imposing a standard simplification.

Assumption 2 (Quasi-Linear Preferences)**.** *Preferences are quasi-linear in effective consumption: $u(\tilde{c}) = \tilde{c}$. Equivalently, $u'(\tilde{c}) = 1$ everywhere, so income effects on labor supply vanish.*

Under Assumption 2, the social welfare weight simplifies to $g(z) = m(a, c(z))/\lambda$. Since welfare is linear in effective consumption, the budget constraint implies the normalization $\lambda = \int_\Theta m(a, c(z))\, h(z)\, dz$, so that $\int_\Theta g(z)\, h(z)\, dz = 1$.

We define two key objects. The *local hazard ratio* at earnings z is

$$\alpha_+(z) \equiv \frac{z\, h(z)}{1 - H(z)},$$

the ratio of the density at z to the fraction of the population earning above z, scaled by z. The *cumulative welfare weight above z* is

$$\bar{g}_+(z) \equiv \frac{\int_z^\infty g(w)\, h(w)\, dw}{1 - H(z)},$$

the population-weighted average social welfare weight among all agents earning above z.

Theorem 2 (Full-Distribution Optimal Marginal Tax Rate)**.** *Under Assumptions 1 and 2, and assuming*

no bunching, the optimal marginal tax rate at any earnings level z in the interior of the support satisfies

$$\frac{T'(z)}{1-T'(z)} = \frac{1-\bar{g}_+(z)}{e_s(z)\,\alpha_+(z)}. \tag{12}$$

Proof. Fix z^* in the interior of the support and consider raising T' by $d\tau$ on the thin interval $[z^*, z^*+\varepsilon]$ for small $\varepsilon > 0$.

Mechanical revenue gain. Every agent earning above z^* now faces a marginal rate that is $d\tau$ higher on the income slice $[z^*, z^*+\varepsilon]$. The revenue gain from this slice is approximately $d\tau \cdot \varepsilon$ per agent above z^*, so

$$dR_{\text{mech}} = d\tau \cdot \varepsilon \cdot [1 - H(z^*)].$$

Mechanical welfare cost. Under Assumption 2, $u' = 1$, so the welfare cost to an agent at earnings $w > z^*$ from the reduction in effective consumption is $m(a, c(w)) \cdot d\tau \cdot \varepsilon$. Integrating over all agents above z^* and using (8):

$$dW_{\text{mech}} = -d\tau \cdot \varepsilon \int_{z^*}^{\infty} m(a, c(w))\, h(w)\, dw = -\lambda\, d\tau \cdot \varepsilon \int_{z^*}^{\infty} g(w)\, h(w)\, dw.$$

Behavioral revenue loss. Agents with earnings in $[z^*, z^*+\varepsilon]$ face a higher marginal tax on this interval. Their effective net-of-tax rate falls from $s(z^*) = n(z^*)\, m(a, c(z^*))$ by $ds = m(a, c(z^*)) \cdot (-d\tau)$. By definition of $e_s(z^*)$:

$$dz = e_s(z^*) \cdot \frac{ds}{s(z^*)} \cdot z^* = -e_s(z^*) \cdot \frac{d\tau}{n(z^*)} \cdot z^*.$$

Each agent in this interval who reduces earnings loses $T'(z^*) \cdot |dz|$ in tax revenue to the government. With density $h(z^*)$ at z^* and interval width ε:

$$dR_{\text{beh}} = -\frac{T'(z^*)}{1-T'(z^*)} \cdot e_s(z^*) \cdot z^* \cdot h(z^*) \cdot \varepsilon \cdot d\tau.$$

Optimality. Setting $dW_{\text{mech}} + \lambda(dR_{\text{mech}} + dR_{\text{beh}}) = 0$ and dividing by $\lambda\, d\tau\, \varepsilon$:

$$-\int_{z^*}^{\infty} g(w)\, h(w)\, dw + [1 - H(z^*)] - \frac{T'(z^*)}{1-T'(z^*)} \cdot e_s(z^*) \cdot z^* \cdot h(z^*) = 0.$$

Dividing through by $1 - H(z^*)$ and recognizing the definitions of $\bar{g}_+(z^*)$ and $\alpha_+(z^*)$:

$$-\bar{g}_+(z^*) + 1 - \frac{T'(z^*)}{1-T'(z^*)} \cdot e_s(z^*) \cdot \alpha_+(z^*) = 0.$$

Rearranging yields (12). □

Remark 3 (Nesting and boundary conditions)**.** *Theorem 2 nests Theorem 1 as a limiting case. In the Pareto tail where $h(z) \propto z^{-(\alpha+1)}$, the local hazard ratio $\alpha_+(z)$ converges to the Pareto parameter α, and $\bar{g}_+(z)$ converges to the income-weighted average weight $\bar{g}$ of Theorem 1. The full-distribution formula thus recovers the top-bracket result in the tail. At the lower end of the support, if $H(0) = 0$, then $\bar{g}_+(0) = \int g(z)\, h(z)\, dz = 1$ by the budget normalization, so $T'(0)/(1-T'(0)) = 0$, meaning the optimum imposes no marginal distortion at the very bottom of the earnings distribution. This is the standard no-distortion-at-the-bottom result.*

Remark 4 (Shape of the optimal rate schedule under frictions)**.** *The friction-adjusted rate schedule inherits the U-shaped pattern of Diamond [1998], with systematically lower rates throughout the interior of the distribution when $\pi_c > 0$. The reason is that $\pi_c > 0$ implies $m_c(a, c) > 0$ (marginal effectiveness rises with income), which means the friction-adjusted welfare weight $g(z) = m(a, c(z))/\lambda$ is increasing in z. An increasing $g(z)$ compresses the redistribution motive: the cumulative average $\bar{g}_+(z)$ is pushed closer to one, shrinking the numerator $1-\bar{g}_+(z)$ in (12). Simultaneously, when m is increasing in c, the composite elasticity $e_s(z)$ exceeds the partial elasticity $e_n(z)$ (see Remark 5 below), enlarging the denominator. Both effects push the optimal friction-adjusted rate below its frictionless counterpart at every income level in the interior of the distribution.*

The following corollary extends the top-bracket bias result to the full distribution.

Corollary 1 (Full-Distribution Over-Taxation Bias)**.** *Under Assumptions 1 and 2, suppose $e_s(z) > e_n(z)$ at every z in the interior of the support. Let T^{bench} be the tax schedule obtained by applying formula (12) with $e_n(z)$ in place of $e_s(z)$, holding $\bar{g}_+(z)$ and $\alpha_+(z)$ fixed. Then $T^{\text{bench}\prime}(z) > T^{*\prime}(z)$ at every interior earnings level: benchmark computations that abstract from access frictions overstate optimal marginal rates throughout the income distribution.*

Proof. At each earnings level z, the benchmark computation inserts $e_n(z) < e_s(z)$ into the denominator of the right-hand side of (12), while holding the numerator $1 - \bar{g}_+(z)$ and the hazard ratio $\alpha_+(z)$ fixed. A smaller denominator yields a larger ratio $T'/(1-T')$, and since the map $\tau \mapsto \tau/(1-\tau)$ is strictly increasing on $[0, 1)$, a larger ratio implies a larger optimal marginal rate. Hence $T^{\text{bench}\prime}(z) > T^{*\prime}(z)$ at every interior z. □

8 Attenuation of Income Redistribution

Before turning to the bias comparison, we formalize a basic intuition about what access frictions do to the transmission from disposable income to welfare. The object $m(a, c)$ measures this transmission at the margin. The following result shows that a worsening of the access environment (a decline in π and a non-increase in π_c) reduces m and hence weakens the marginal effectiveness of disposable income as a vehicle for welfare.

Proposition 1 (Attenuation of Marginal Effectiveness)**.** *For any (a, c) with $c \geq \underline{c}$, marginal effectiveness is*

$$m(a, c) = \pi(a, c) + \pi_c(a, c)\,(c - \underline{c}). \tag{13}$$

If the access technology deteriorates in the sense that $\pi(a, c)$ decreases pointwise and $\pi_c(a, c)$ does not increase pointwise, then $m(a, c)$ weakly decreases at every (a, c) with $c \geq \underline{c}$.

Proof. We differentiate $\tilde{c} = \underline{c} + \pi(a, c)(c - \underline{c})$ directly with respect to c:

$$\frac{\partial \tilde{c}}{\partial c} = \pi(a, c) + \pi_c(a, c)\,(c - \underline{c}),$$

which establishes (13). Both terms in this expression are non-negative: $\pi > 0$ by Assumption 1, and $\pi_c(c - \underline{c}) \geq 0$ since $\pi_c \geq 0$ and $c \geq \underline{c}$. Under the stated deterioration, $\pi(a, c)$ falls at every point, which directly reduces the first term. The second term $\pi_c(a, c)(c - \underline{c})$ does not increase, since π_c is assumed not to increase and $(c - \underline{c}) \geq 0$ is unchanged. Hence $m(a, c)$ weakly decreases at every (a, c). □

Proposition 1 has a direct welfare interpretation. A deterioration in the access environment attenuates the marginal effectiveness of redistribution. A dollar transferred from high earners to low earners generates less effective consumption at the destination when access frictions are severe. This is distinct from the behavioral distortion. It is a pure transmission loss on the redistribution side of the ledger, and it operates independently of whether the transfer comes through the tax schedule or through access policy directly.

9 The Over-Taxation Bias

We now formalize and sharpen the bias result previewed in Remark 4. The central question is what happens to inferred optimal tax rates when a researcher uses the standard frictionless formula (which measures the earnings elasticity with respect to the net-of-tax rate n) in a world where the true behavioral sufficient statistic is the elasticity with respect to the composite shifter $s = n \cdot m$.

Theorem 3 (Over-Taxation Bias)**.** *Fix a threshold $\bar{y}$ and suppose the equilibrium satisfies $e_s(\bar{y}) > e_n(\bar{y})$. Let $\tau^{\text{bench}}(\bar{y})$ denote the top marginal rate implied by the frictionless Saez formula,*

$$\frac{\tau^{\text{bench}}(\bar{y})}{1-\tau^{\text{bench}}(\bar{y})} = \frac{1-\bar{g}(\bar{y})}{\alpha(\bar{y})\, e_n(\bar{y})},$$

where the objects $\bar{g}(\bar{y})$ and $\alpha(\bar{y})$ are held fixed at their equilibrium values. Let $\tau^*(\bar{y})$ denote the optimal rate from Theorem 1. Then $\tau^{\text{bench}}(\bar{y}) > \tau^*(\bar{y})$.

Proof. Because $e_s(\bar{y}) > e_n(\bar{y}) > 0$, we have

$$\frac{1-\bar{g}(\bar{y})}{\alpha(\bar{y})\, e_s(\bar{y})} < \frac{1-\bar{g}(\bar{y})}{\alpha(\bar{y})\, e_n(\bar{y})}.$$

(Here we use $1-\bar{g}(\bar{y}) > 0$, which holds when the average welfare weight below the threshold exceeds that above it, as is typical for redistributive preferences, and $\alpha(\bar{y}) > 0$.) The map $x \mapsto x/(1+x)$ is strictly increasing on $\mathbb{R}_+$, so a smaller value of the ratio $\tau/(1-\tau)$ implies a strictly smaller value of τ. It follows that the rate satisfying the frictionless formula, τ^{bench}, strictly exceeds the rate satisfying the friction-corrected formula, τ^*. □

Remark 5 (When does $e_s > e_n$ hold?)**.** *The composite elasticity e_s captures the total response of earnings to a change in the effective return to work, allowing $m(a,c)$ to adjust as disposable income changes. The partial elasticity e_n holds m fixed. Whenever $m(a,c)$ is increasing in c (which follows from $\pi_c > 0$ and mild conditions on π_{cc}), a reduction in the net-of-tax rate n that causes earnings (and thus c) to fall also causes m to fall, amplifying the incentive effect. Formally, for a small reduction $dn < 0$:*

$$d\log s = d\log n + d\log m = d\log n + m_c\, dc/m \approx d\log n\,\big(1 + m_c\, c/m\big),$$

so $e_s = e_n \cdot (1+\gamma) > e_n$ when $\gamma \equiv m_c\, c/m > 0$. This is not merely a technical assumption. It is the natural implication of the empirically documented income gradient in access frictions. Section 12 quantifies γ from experimental and survey evidence and reports the resulting magnitude of the bias.

10 Access Frictions and the Extensive Margin

Sections 6–9 characterize optimal marginal taxation of agents who are already in the labor market. A central insight of Saez [2002] is that, for low-income households, the relevant margin of response to tax incentives is often whether to work at all rather than how much to work. We now embed an explicit participation decision into the model and examine how access frictions modify the optimal participation tax formula.

10.1 Model with Participation

Augment the baseline model with a heterogeneous fixed cost of labor-market participation. An agent of type (θ, q) incurs cost $q \geq 0$ upon entering the labor market, where q is drawn independently of θ from a distribution with CDF Ψ and density ψ, which are continuous and positive on $\mathbb{R}_+$. The cost q represents the heterogeneous personal costs of working: childcare, transportation, foregone domestic production, and other idiosyncratic barriers that vary across households but are not directly observable by the government.

Non-employment generates disposable income $c_0 \equiv -T(0)$ (the government may pay a lump-sum transfer to non-workers) and effective consumption $\tilde{c}_0 \equiv \underline{c} + \pi(a, c_0)\,(c_0 - \underline{c})$. An agent of type (θ, q) participates in the labor market if and only if the net utility gain from doing so is non-negative:

$$u(\tilde{c}(y(\theta))) - v(y(\theta)/\theta) - u(\tilde{c}_0) \;\geq\; q. \tag{14}$$

All other elements of the model (the access technology, the government's objective, the budget constraint) remain as before.

Notice that access frictions enter condition (14) on both sides. They affect the effective consumption from employment through $\tilde{c}(y(\theta))$ and from non-employment through $\tilde{c}_0$. The net effect depends on the relative severity of frictions in the two states. When frictions are more severe for employed households (for instance, because accessing formal financial services or navigating employer benefit systems is costly), participation is less attractive than the statutory wage gain suggests. When frictions are more severe for non-employed households (because accessing welfare programs requires navigating complex administrative processes), the utility gain from work is amplified relative to the statutory income gain.

Definition 1 (Effective Surplus and Participation Elasticity)**.** *For an agent of type θ, the* effective participation surplus *is the utility gain from working:*

$$\tilde{S}(\theta) \;\equiv\; u(\tilde{c}(y(\theta))) - v(y(\theta)/\theta) - u(\tilde{c}_0).$$

An agent participates if and only if $q \leq \tilde{S}(\theta)$, so the participation rate for type θ is $P(\theta) = \Psi(\tilde{S}(\theta))$. The friction-adjusted participation elasticity *with respect to employment disposable income $c(y(\theta))$ is*

$$\eta^c_P(\theta) \;\equiv\; \frac{d\log P(\theta)}{d\log c(y(\theta))} = \psi(\tilde{S}(\theta)) \cdot u'(\tilde{c}(y(\theta))) \cdot m(a, c(y(\theta))) \cdot \frac{c(y(\theta))}{P(\theta)}.$$

The participation elasticity η^c_P depends on both $u'(\tilde{c}) \cdot m$ and the density ψ of participation costs near the margin. The factor $u'(\tilde{c}) \cdot m$ is the marginal utility of effective consumption from employment (and hence the marginal increase in $\tilde{S}$ per dollar of additional disposable income), so η^c_P measures the behavioral response of the participation margin to a dollar of income change, fully accounting for the friction wedge.

10.2 Optimal Participation Tax

Proposition 2 (Optimal Participation Tax under Access Frictions)**.** *Under Assumption 1, the optimal participation tax*

$$\rho^*(\theta) \equiv \frac{T(y(\theta)) - T(0)}{c(y(\theta)) - c_0}$$

satisfies

$$\rho^*(\theta) = \frac{1 - g_1(\theta)}{\eta_P^c(\theta)}, \tag{15}$$

where $g_1(\theta) \equiv u'(\tilde{c}(y(\theta))) \cdot m(a, c(y(\theta)))/\lambda$ is the friction-adjusted welfare weight of a participating agent of type θ.

Proof. Fix a type θ and consider a small increase $d\tau$ in $T(y(\theta))$ holding $T(0)$ fixed, that is, a marginal increase in the participation tax on type-θ workers. This reform has three effects.

Mechanical welfare cost. Among the $P(\theta)$ fraction of type-θ agents who currently participate, each loses $m(a, c(y(\theta))) \cdot d\tau$ units of effective consumption (since $\partial\tilde{c}/\partial c = m$ and $dc = -d\tau$). The total welfare cost is

$$dW_{\text{mech}} = -u'(\tilde{c}(y(\theta))) \cdot m(a, c(y(\theta))) \cdot P(\theta) \cdot d\tau = -\lambda\, g_1(\theta) \cdot P(\theta) \cdot d\tau.$$

Mechanical revenue gain. The $P(\theta)$ participating agents each pay $d\tau$ more:

$$dR_{\text{mech}} = P(\theta) \cdot d\tau.$$

Participation response. The reform reduces the effective surplus $\tilde{S}(\theta)$ by $u'(\tilde{c}(y(\theta))) \cdot m(a, c(y(\theta))) \cdot d\tau$. The marginal agents (those with q exactly at $\tilde{S}(\theta)$) now find it optimal not to work. The change in the participation rate is

$$dP(\theta) = -\psi(\tilde{S}(\theta)) \cdot u'(\tilde{c}) \cdot m \cdot d\tau.$$

Each agent who leaves the labor market costs the government the tax they were paying, $T(y(\theta)) - T(0) = \rho^* \cdot (c(y(\theta)) - c_0)$. The net revenue effect from the participation margin is

$$dR_{\text{part}} = [T(y(\theta)) - T(0)] \cdot dP(\theta) = -\rho^*(\theta) \cdot \eta_P^c(\theta) \cdot P(\theta) \cdot d\tau,$$

where the last equality uses the definition of η_P^c from Definition 1.

Optimality. Setting the total first-order effect on the Lagrangian to zero:

$$dW_{\text{mech}} + \lambda\big(dR_{\text{mech}} + dR_{\text{part}}\big) = 0.$$

Dividing by $\lambda\, P(\theta)\, d\tau$:

$$-g_1(\theta) + 1 - \rho^*(\theta)\, \eta_P^c(\theta) = 0.$$

Solving for $\rho^*(\theta)$ yields (15). □

Remark 6 (Economic content of the friction-adjusted participation formula)**.** *Equation (15) is formally identical to the Saez [2002] participation tax condition, with two friction-induced modifications. First, $g_1(\theta)$ incorporates $m(a, c(y(\theta)))$: the social value of a marginal dollar delivered to an employed agent is rescaled by how much of that dollar converts into effective consumption. If employed agents face heavy frictions (m low),*

they are worth less in welfare terms per dollar of statutory income, and the case for subsidizing employment through a negative ρ^* is correspondingly weakened. Second, η_P^c is proportional to $u'(\tilde{c}) \cdot m$ rather than $u'(c)$: the participation response to a statutory income change is attenuated by the extent to which the additional statutory dollar actually reaches the agent as consumption. In the frictionless benchmark, $m = 1$ and $\tilde{c} = c$, and condition (15) reduces exactly to the Saez (2002) formula.

Corollary 2 (Direction of Participation Tax Bias)**.** *Let $\rho^{\text{bench}}(\theta)$ denote the benchmark participation tax computed using the frictionless formula (setting $m \equiv 1$ and $\tilde{c} = c$ throughout). Under Assumption 2:*

(i) If $m(a, c(y(\theta))) < 1$ at the participation margin (so that frictions are heavier for employed agents), then $\eta_P^c < \eta_P^{0,c}$ (participation responds less to statutory incentives), and the optimal participation tax ρ^* is higher (less negative) than the benchmark predicts. Standard EITC-style analyses that ignore employment frictions will prescribe excessive in-work subsidies.

(ii) If $m(a, c(y(\theta))) > 1$ (which can occur when non-employment frictions are more severe than employment frictions, so that the access intensity π rises sharply with income near the participation margin), then $\eta_P^c > \eta_P^{0,c}$, and the optimal participation tax is lower (more negative) than the benchmark predicts.

Proof. Under quasi-linearity $u' = 1$, the friction-adjusted participation elasticity is $\eta_P^c = \psi(\tilde{S}) \cdot m \cdot c(y(\theta))/P(\theta)$ and its frictionless counterpart is $\eta_P^{0,c} = \psi(\tilde{S}^0) \cdot c(y(\theta))/P^0(\theta)$. When $m < 1$ and participation costs are similarly distributed in both cases ($\psi(\tilde{S}) \approx \psi^0(\tilde{S}^0)$ and $P \approx P^0$ to first order), we have $\eta_P^c < \eta_P^{0,c}$, which makes the denominator of (15) smaller relative to the benchmark. With the numerator $(1 - g_1)$ approximately unchanged (since $g_1 \approx m/\lambda < 1/\lambda = g_1^0$ when $m < 1$, though this numerator effect is dominated by the denominator effect under the stated approximation), the optimal rate ρ^* is higher. The case $m > 1$ follows by the same logic with reversed inequalities. □

11 Joint Optimal Policy

The results so far have analyzed the optimal tax schedule taking the access policy level a as given. We now ask how the government should simultaneously choose $T(\cdot)$ and a. This joint optimization problem captures the idea that redistribution has two levers (the statutory schedule and the institutional infrastructure) and that the two instruments interact through welfare weights and effective elasticities.

11.1 Optimal Access Conditional on the Tax Schedule

We first characterize the optimal access level for a fixed tax schedule, which establishes the key first-order condition that will also appear in the joint optimum.

Proposition 3 (Optimal Access for a Fixed Schedule)**.** *Fix a tax schedule $T(\cdot)$. At any interior optimum, the optimal access level a^* satisfies*

$$\kappa = \int_{\Theta} u'(\tilde{c}(\theta)) \cdot \pi_a(a, c(\theta)) \cdot (c(\theta) - \underline{c})\, f(\theta)\, d\theta. \tag{16}$$

Proof. We differentiate the Lagrangian $\mathcal{L} = W(T, a) + \lambda\big[\int T(y(\theta)) f(\theta) d\theta - \kappa a\big]$ with respect to a, holding $T(\cdot)$ fixed so that $c(\theta) = y(\theta) - T(y(\theta))$ is unchanged. Effective consumption depends on a through $\tilde{c}(\theta) = \underline{c} + \pi(a, c(\theta))(c(\theta) - \underline{c})$. Differentiating with respect to a:

$$\frac{\partial \tilde{c}(\theta)}{\partial a} = \pi_a(a, c(\theta))\,(c(\theta) - \underline{c}).$$

The direct budget cost of access policy contributes $-\lambda\kappa$ to $\partial\mathcal{L}/\partial a$. Setting the derivative to zero and normalizing $\lambda = 1$ at the optimum yields (16). □

Equation (16) has a natural interpretation. The optimal level of access investment equates its marginal resource cost κ to its aggregate social benefit, the population-average increase in welfare from a marginal improvement in access quality. Agents with higher disposable income above the floor ($c(\theta) - \underline{c}$) benefit more in absolute terms from a given improvement in π_a, and the welfare weight $u'(\tilde{c}(\theta))$ discounts this by the marginal utility of consumption. The interplay between these two forces determines the optimal targeting of access investment.

11.2 Complementarity Between Access and the Tax Schedule

Before characterizing the full joint optimum, it is useful to isolate a partial-equilibrium complementarity between the two policy instruments.

Proposition 4 (Complementarity)**.** *Suppose $\pi_{ac}(a, c) \geq 0$ for all (a, c) with $c \geq \underline{c}$. Holding $u'(\tilde{c}(y))$, λ, $\alpha(\bar{y})$, and $e_s(\bar{y})$ fixed, an increase in the access level a weakly increases every welfare weight $g(z)$ and hence the average weight $\bar{g}(\bar{y})$, implying that the optimal top marginal tax rate weakly decreases.*

Proof. Differentiating $m(a, c) = \pi(a, c) + \pi_c(a, c)(c - \underline{c})$ with respect to a:

$$m_a(a, c) = \pi_a(a, c) + \pi_{ac}(a, c)\,(c - \underline{c}).$$

Under the hypothesis, $\pi_a > 0$ (Assumption 1), $\pi_{ac} \geq 0$, and $c \geq \underline{c}$, so $m_a(a, c) \geq 0$ at every point. Since $g(z) = u'(\tilde{c}(z)) \cdot m(a, c(z))/\lambda$ and we treat $u'(\tilde{c}(z))$ and λ as fixed, $g(z)$ weakly increases pointwise in a. The income-weighted average $\bar{g}(\bar{y})$ therefore weakly increases. From formula (11), $\tau(\bar{y})$ is strictly decreasing in $\bar{g}(\bar{y})$ holding α and e_s fixed. □

Remark 7. *Proposition 4 identifies a channel through which better access policy enables lower marginal tax rates. When a rises, each agent is better able to translate disposable income into effective consumption, which raises welfare weights throughout the distribution and compresses the redistributive motive measured in welfare units. This is a partial-equilibrium result: it holds the earnings distribution, the composite elasticity, and the shadow cost of public funds fixed at their pre-reform values. In general equilibrium, the earnings distribution and the shadow cost would adjust, potentially reinforcing or moderating this complementarity. The result should therefore be read as isolating the welfare-weight channel rather than as a global characterization of the joint optimum.*

11.3 Characterizing the Joint Optimum

We now characterize the pair (T^*, a^*) that simultaneously maximizes social welfare subject to incentive compatibility and the budget constraint.

Theorem 4 (Joint Optimum)**.** *Suppose the government's Lagrangian is jointly concave in (T, a) and that, for each fixed T, the first-order condition (16) has a unique solution in a. The jointly optimal policy (T^*, a^*) is the unique fixed point of the system:*

(i) T^* *satisfies the full-distribution optimality condition of Theorem 2 (under quasi-linearity, or Theorem 1 in the top bracket) with welfare weights evaluated at* a^**:*

$$g^*(z) = \frac{u'(\tilde{c}^*(z)) \cdot m(a^*, c^*(z))}{\lambda^*};$$

(ii) a^* *satisfies condition (16) evaluated at the income distribution induced by* T^**:*

$$\kappa = \int_\Theta u'(\tilde{c}^*(\theta)) \cdot \pi_a(a^*, c^*(\theta)) \cdot (c^*(\theta) - \underline{c})\, f(\theta)\, d\theta,$$

where $c^*(\theta) = y^*(\theta) - T^*(y^*(\theta))$.

At the joint optimum, the budget constraint holds with equality, and the marginal social value of access investment equals its marginal cost κ.

Proof. We characterize the optimum through necessary conditions on the Lagrangian

$$\mathcal{L}(T, a, \lambda) = W(T, a) + \lambda \left[\int_\Theta T(y(\theta)) f(\theta)\, d\theta - \kappa a \right].$$

The first-order condition with respect to variations in T at each income level reproduces the perturbation arguments of Theorems 1 and 2, with a held at its optimized value a^*. These conditions yield the optimal schedule described in part (i). The first-order condition with respect to a, derived in Proposition 3, yields condition (ii). Joint concavity of the Lagrangian ensures that the system of first-order conditions is necessary and sufficient for a global maximum.

To show uniqueness of the fixed point, define the operator $\mathcal{T}$ that maps a policy pair $(T^{(k)}, a^{(k)})$ to the updated pair $(T^{(k+1)}, a^{(k+1)})$ by solving conditions (i) and (ii) at each iteration. Joint strict concavity of $\mathcal{L}$ implies that $T^{(k)}$ converges to the unique optimal schedule for any fixed $a^{(k)}$, and $a^{(k+1)}$ is the unique solution to (16) given $T^{(k)}$ by assumption. The joint solution is the unique fixed point of $\mathcal{T}$, and it exists by the Schauder fixed-point theorem applied to the compact feasible set defined by the budget constraint. □

Remark 8 (The coupling between the two instruments)**.** *The two conditions in Theorem 4 are deeply coupled. The optimal tax schedule depends on welfare weights* $g^*(z)$ *that incorporate* $m(a^*, c^*(z))$*, and the optimal access level depends on the income distribution* $c^*(\theta)$ *induced by* T^**. A change in the access level alters welfare weights and therefore shifts the optimal schedule. A change in the tax schedule alters the income distribution and therefore shifts the optimal access level. The partial complementarity established in Proposition 4 captures the first direction. The second direction (how tax policy feeds back into optimal access investment) is characterized by Proposition 5 below. Together, these results confirm that treating access policy and tax design as separable problems, as is standard in both the optimal tax literature and the program-design literature, will in general yield suboptimal policies.*

11.4 Inequality and the Optimal Level of Access Investment

The most policy-relevant comparative static concerns how the optimal access investment a^* responds to changes in income inequality. When income is more concentrated at the bottom of the distribution, households face more severe access constraints (since m is lower at low c), the social value of improving access is higher, and the optimal investment a^* should rise. The following proposition makes this precise.

Proposition 5 (Inequality and Access Investment)**.** *Define the* aggregate social return to access investment *as*

$$\Gamma(a, c(\cdot)) \equiv \int_{\Theta} u'(\tilde{c}(\theta)) \cdot \pi_a(a, c(\theta)) \cdot (c(\theta) - \underline{c})\, f(\theta)\, d\theta.$$

Suppose the distribution of disposable incomes $c(\theta)$ shifts downward in the sense of first-order stochastic dominance (FOSD). Then Γ increases (and the optimal access investment a^* increases) if and only if the integrand $\phi(c) \equiv u'(\tilde{c}) \cdot \pi_a(a, c) \cdot (c - \underline{c})$ is non-increasing in c.

A sufficient condition for ϕ to be non-increasing is that the coefficient of absolute risk aversion $\mathcal{A}(c) \equiv -u''(\tilde{c}) \cdot m(a, c)/u'(\tilde{c})$ satisfies, for all $c \geq \underline{c}$:

$$\mathcal{A}(c) \;\geq\; \frac{\pi_{ac}(a, c)\,(c - \underline{c}) + \pi_a(a, c)}{m(a, c)\,\pi_a(a, c)\,(c - \underline{c})}. \tag{17}$$

Proof. **FOSD direction.** Since $\Gamma = \int \phi(c(\theta))\, f(\theta)\, d\theta = \mathbb{E}[\phi(c)]$, a downward FOSD shift increases $\mathbb{E}[\phi(c)]$ if and only if ϕ is non-increasing. A FOSD shift toward lower c values puts more weight on regions where ϕ is evaluated, and $\mathbb{E}[\phi]$ rises if and only if ϕ falls as c rises.

Sufficient condition. We compute $\phi'(c)$ directly. Since $\tilde{c} = \underline{c} + \pi(a, c)(c - \underline{c})$, we have $d\tilde{c}/dc = m(a, c)$. Applying the product rule:

$$\phi'(c) = u''(\tilde{c}) \cdot m(a, c) \cdot \pi_a(a, c) \cdot (c - \underline{c}) + u'(\tilde{c}) \cdot \big[\pi_{ac}(a, c)\,(c - \underline{c}) + \pi_a(a, c)\big].$$

Dividing through by $u'(\tilde{c}) \cdot \pi_a(a, c) \cdot (c - \underline{c}) > 0$:

$$\frac{\phi'(c)}{u'(\tilde{c}) \cdot \pi_a(a, c) \cdot (c - \underline{c})} = \frac{u''(\tilde{c}) \cdot m(a, c)}{u'(\tilde{c})} + \frac{\pi_{ac}(a, c)\,(c - \underline{c}) + \pi_a(a, c)}{\pi_a(a, c)\,(c - \underline{c})}.$$

The condition $\phi'(c) \leq 0$ is equivalent to the first term being sufficiently negative to dominate the second, which is exactly condition (17), upon noting that $u''(\tilde{c}) \cdot m/u'(\tilde{c}) = -\mathcal{A}(c)$.

Comparative static on a^*. At an interior solution, $\Gamma(a^*, c(\cdot)) = \kappa$. A downward FOSD shift increases Γ (when ϕ is non-increasing). To restore equality, a^* must rise. To see that a higher a^* reduces Γ, note that $\partial\Gamma/\partial a = \int u'(\tilde{c})\, \pi_{aa}\, (c - \underline{c})\, f\, d\theta + \int u''(\tilde{c})\, m\, [\pi_a(c - \underline{c})]^2\, f\, d\theta$, which is negative when $\pi_{aa} \leq 0$ and $u'' < 0$ (strict concavity and non-increasing returns to access investment). The implicit function theorem then gives $da^*/d[\text{FOSD decrease}] > 0$. □

Remark 9 (On condition (17) and its empirical plausibility)**.** *Condition (17) requires effective absolute risk aversion to exceed a threshold determined by the curvature of the access technology. In the simplest case where $\pi_{ac} = 0$ and π_a is independent of c, the right-hand side collapses to $1/(c - \underline{c})$. The condition then says $\mathcal{A}(c) \geq 1/(c - \underline{c})$, which requires that effective risk aversion exceed the inverse of the consumption gap above the outside option. For CRRA utility $u(\tilde{c}) = \tilde{c}^{1-\sigma}/(1 - \sigma)$ with curvature parameter σ, one can verify that $\mathcal{A}(c) = \sigma \cdot m/\tilde{c}$, and condition (17) holds at c reasonably above $\underline{c}$ whenever $\sigma \geq 1$, a standard assumption in quantitative work. The condition is therefore not restrictive in applications.*

Proposition 5 has a direct policy implication: more unequal societies should not only tax more (the standard Mirrlees lesson) but also invest more in access infrastructure (better service delivery, lower administrative burdens, and improved program navigation), because the social return to each unit of access improvement is higher when effective consumption at the bottom of the distribution is more scarce.

12 Calibrating the Bias

The theoretical results establish that benchmark computations using the statutory net-of-tax elasticity e_n overstate the optimal marginal tax rate whenever the friction gradient $\gamma \equiv m_c \cdot c/m$ is positive. This section asks how large the correction is in practice. We calibrate γ from available empirical evidence, pair it with standard estimates of the remaining sufficient statistics, and compute the implied friction-corrected optimal top rate τ^* alongside the benchmark τ^{bench}.

12.1 Calibration Strategy

The friction-corrected formula (Theorem 1) is

$$\frac{\tau^*}{1-\tau^*} = \frac{1-\bar{g}}{\alpha \cdot e_s}, \qquad e_s = e_n \cdot (1+\gamma). \tag{18}$$

The benchmark sets $\gamma = 0$, using e_n in place of e_s. The correction $\tau^{\text{bench}} - \tau^*$ is determined by four objects: the elasticity e_n, the Pareto tail parameter α, the average welfare weight $\bar{g}$, and the friction gradient γ. We calibrate each in turn.

Earnings elasticity e_n. We use three values that span the empirical range. Diamond and Saez [2011] and Saez [2001] adopt $e_n = 0.25$ as a conservative benchmark, which implies a revenue-maximizing rate of about 73 %. Gruber and Saez [2002] estimate $e_n = 0.57$ for upper income earners using variation in U.S. marginal rates. A mid-range value of $e_n = 0.40$, consistent with the central range surveyed by Saez et al. [2012], provides an intermediate case.

Pareto tail parameter α. We set $\alpha = 1.5$, the value used by Diamond [1998] and Diamond and Saez [2011] for the U.S. earnings distribution. Atkinson et al. [2011] document that income tail indices range from 1.5 to 3 across countries and time periods, with the U.S. top-income distribution in recent decades closer to the lower end of this range. We verify robustness at $\alpha = 2.0$ in Panel C.

Welfare weight $\bar{g}$. Following the standard practice in the sufficient-statistics literature [Diamond and Saez, 2011], we report results for $\bar{g} = 0$ (the revenue-maximizing case, in which the government places zero social weight on top earners' income) and $\bar{g} = 0.20$ (a moderate redistributive case).

Friction gradient γ. The friction gradient $\gamma = m_c \cdot c/m$ measures the elasticity of marginal effectiveness with respect to disposable income. For a parametric access technology of the form $\pi(a, c) \propto c^{\beta}$, differentiation shows $\gamma = \beta$ throughout the income distribution. The parameter β is thus the income elasticity of the household's ability to convert disposable income into effective consumption.

Experimental evidence from SNAP. Finkelstein and Notowidigdo [2019] conduct a field experiment with 30,000 elderly SNAP-eligible non-participants. Information plus application assistance raises enrollment from 6 % to 18 %, implying that roughly 82 % of eligible individuals face access frictions sufficient to deter take-up even when the financial gain is substantial. The application cost per complier, relative to the average SNAP benefit of approximately $1,700 per year, implies a friction cost of 2 to 4 per cent of annual income for these near-poor households. Critically, the authors find that those induced to enroll by the intervention have *higher* net incomes and are *less* sick than the average existing enrollee, confirming that frictions are more severe precisely for households with the fewest resources.

Income gradient of administrative burdens. Herd and Moynihan [2018] document that compliance costs, learning costs, and psychological costs of navigating transfer programs are substantially higher for lower-income households as a fraction of income. Their framework implies that the friction burden falls by roughly 10 to 20 per cent (relative) for each doubling of income in the range below $100,000. Combining with estimates from Currie [2006], who finds that low-income families leave approximately $600 to $1,200 per year in unclaimed benefits, and evidence on differential financial services access (the premium on informal versus formal credit is worth roughly 3 to 6 per cent of annual income for the bottom income quintile), the total income gradient of effective consumption access corresponds to $\gamma \approx 0.10$ to 0.20 for households in the bottom two quintiles.

Top earners. For top earners, frictions are minimal and the income gradient is flat. Administrative compliance costs as a fraction of income are negligible for high-income households who employ tax professionals and financial advisors. We therefore adopt $\gamma = 0.05$ as the friction gradient at the top, consistent with evidence that high-income households face de minimis take-up frictions for any programs they are eligible for.

These considerations motivate three calibration scenarios:

- *Low:* $\gamma = 0.05$. Applicable to top earners in the U.S., where access frictions are low and the income gradient of access is flat. Provides a lower bound on the bias correction.
- *Moderate:* $\gamma = 0.10$. Applicable to the broad U.S. working population, integrating the gradient over the middle of the income distribution. Consistent with the combined evidence from Finkelstein and Notowidigdo and Herd and Moynihan.
- *High:* $\gamma = 0.20$. Applicable to environments with steeper access frictions, such as developing-country tax-transfer systems or U.S. programs with intensive documentation requirements. Consistent with evidence on program leakage and informality premiums in lower-income settings [Hernanz et al., 2004, Ko and Moffitt, 2022].

12.2 Results

Table 1 reports the friction-corrected optimal top marginal tax rate τ^* and the benchmark τ^{bench} for each combination of parameters, expressed as percentages.

Several patterns emerge. First, the friction correction is consistent across all parameter combinations. Regardless of the assumed elasticity, tail parameter, or welfare weight, incorporating the access-friction gradient reduces the implied optimal top rate. Second, the magnitude of the correction is economically meaningful. Even the low scenario ($\gamma = 0.05$), calibrated conservatively for top earners, shaves roughly 1 percentage point off the optimal rate. The moderate scenario ($\gamma = 0.10$), appropriate for the broad U.S. population, delivers a correction of 1.5 to 4 percentage points, which is on the order of the policy-relevant differences between major tax reform proposals. The high scenario ($\gamma = 0.20$), applicable to programs with steep access gradients, implies a downward correction of 3.5 to 4.5 percentage points. Third, the relative bias is roughly constant across elasticity values within each panel, since $\tau^{\text{bench}} - \tau^* \approx \tau^{\text{bench}} \cdot \gamma/(1+\gamma)$ to a first-order approximation.[2]

[2]From the formula: $\tau^* \approx \tau^{\text{bench}}/(1+\gamma)$ for small γ, so the absolute bias is approximately $\tau^{\text{bench}} \cdot \gamma/(1+\gamma)$.

Table 1: Benchmark and friction-corrected optimal top marginal tax rates

	Benchmark	Friction gradient γ			
Elasticity e_n	$\gamma = 0$	$\gamma = 0.05$	$\gamma = 0.10$	$\gamma = 0.20$	Bias range
Panel A: $\alpha = 1.5$, $\bar{g} = 0$ *(revenue-maximising)*					
$e_n = 0.25$	72.7	71.7	70.8	69.0	1.0–3.8 pp
$e_n = 0.40$	62.5	61.3	60.2	58.1	1.2–4.4 pp
$e_n = 0.57$	53.9	52.7	51.5	49.4	1.2–4.6 pp
Panel B: $\alpha = 1.5$, $\bar{g} = 0.20$ *(redistributive)*					
$e_n = 0.25$	68.1	67.0	66.0	64.0	1.1–4.1 pp
$e_n = 0.40$	57.1	55.9	54.8	52.6	1.2–4.5 pp
$e_n = 0.57$	48.3	47.1	46.0	43.8	1.2–4.5 pp
Panel C: $\alpha = 2.0$, $\bar{g} = 0$ *(thinner tail)*					
$e_n = 0.25$	66.7	65.6	64.5	62.5	1.1–4.2 pp
$e_n = 0.40$	55.6	54.3	53.2	51.0	1.2–4.5 pp
$e_n = 0.57$	46.7	45.5	44.4	42.2	1.2–4.5 pp

Notes. Each cell reports the optimal top marginal tax rate as a percentage, computed from $\tau/(1-\tau) = (1-\bar{g})/[\alpha \cdot e_n(1+\gamma)]$. Bias range gives the friction correction $\tau^{\text{bench}} - \tau^*$ in percentage points across the three non-benchmark columns. Parameters: $e_n \in \{0.25, 0.40, 0.57\}$ from Diamond and Saez [2011], Saez et al. [2012], and Gruber and Saez [2002] respectively; $\alpha \in \{1.5, 2.0\}$ from Diamond [1998] and Atkinson et al. [2011]; $\bar{g} \in \{0, 0.20\}$ following Diamond and Saez [2011]; $\gamma \in \{0.05, 0.10, 0.20\}$ calibrated from Finkelstein and Notowidigdo [2019], Herd and Moynihan [2018], and Ko and Moffitt [2022] as described in the text.

12.3 Distribution-Wide Bias

The full-distribution formula (Theorem 2) implies that the bias from omitting frictions extends to every income level in the interior of the distribution. The magnitude of the correction at each income level z is proportional to $\gamma(z)$, the local friction gradient at that point. Since empirical evidence strongly suggests that frictions are more income-graded at the bottom of the distribution than the top, the downward correction to optimal rates is largest precisely where the U-shaped rate schedule is also highest: at lower income levels, where the redistributive pressure is greatest but the access environment is worst.

Table 2 reports the ratio $\tau^*(z)/\tau^{\text{bench}}(z) = 1/(1+\gamma(z))$ at three income levels.

Table 2: Relative bias correction across the income distribution

Income level	Approx. U.S. percentile	$\gamma(z)$	Correction factor	Rate as % of benchmark
Lower income	P20–P30	0.30	0.769	76.9 %
Middle income	P50–P60	0.15	0.870	87.0 %
Top bracket	P99+	0.05	0.952	95.2 %

Notes. Each row reports the ratio $\tau^*(z)/\tau^{\text{bench}}(z) = 1/(1+\gamma(z))$ at the corresponding income level. Friction gradients $\gamma(z)$ are calibrated from Finkelstein and Notowidigdo [2019] and Herd and Moynihan [2018]: $\gamma = 0.30$ at the bottom (reflecting the steep income gradient of SNAP, Medicaid, and housing assistance take-up), $\gamma = 0.15$ at the median, and $\gamma = 0.05$ at the top. All other sufficient statistics are held constant across rows.

At the bottom of the income distribution, the friction correction reduces the optimal rate to roughly 77 per cent of the frictionless benchmark. For the middle of the distribution the correction is about 13 per

cent, and at the top it is about 5 per cent. Since the frictionless optimal rate schedule is U-shaped, the friction correction disproportionately reduces rates where they are highest in the frictionless benchmark, compressing the U-shape and generating a flatter, uniformly lower friction-adjusted rate schedule.

12.4 Caveats

Three qualifications are worth stating explicitly. First, γ is calibrated from a combination of program-specific evidence and administrative burden estimates rather than from a single clean identification strategy. The estimates should be read as order-of-magnitude calibrations. The key finding is robust: any $\gamma > 0$ produces a downward correction, and even conservative values (0.05 to 0.10) generate corrections of 1 to 4 percentage points.

Second, the calibration holds $\bar{g}$ fixed at its frictionless value. In principle, access frictions also modify welfare weights through $m(a, c)$ as established in Theorem 1 and Proposition 1: the welfare weight at each income level is $g(z) = u'(\tilde{c}(z)) \cdot m(a, c(z))/\lambda$, which is lower than the frictionless weight whenever $m < 1$. Accounting for this second channel would reinforce the downward correction to optimal rates, so the numbers in Tables 1 and 2 understate the full magnitude of the bias.

Third, the model abstracted from avoidance responses and multiple behavioral margins. Piketty et al. [2014] show that accounting for avoidance and bargaining responses substantially lowers optimal top rates relative to the pure labor supply case. Incorporating access frictions into the three-elasticity framework of Piketty et al. would add the friction gradient to each behavioral channel, amplifying the correction identified here.

13 Discussion

13.1 Two Instruments, One Redistribution Problem

The model reveals a fundamental symmetry between the two instruments available to the government. The tax schedule $T(\cdot)$ operates on disposable income: it moves resources from high earners to the rest of the population, subject to the behavioral response that high earners reduce their earnings in response. The access level a operates on the mapping from income to welfare: it determines how efficiently the resources that reach each household translate into effective consumption, subject to the resource cost κa. In the standard model, only the first instrument exists, and the entire redistribution problem reduces to designing $T(\cdot)$ optimally. In the present model, both instruments matter, and they interact.

Theorem 4 shows that the optimal joint policy cannot be decomposed into two separate problems. The welfare weights that govern optimal tax design depend on a through the marginal effectiveness m, and the income distribution that governs optimal access investment depends on T through agents' behavioral responses. A government that designs its tax schedule ignoring the access environment will systematically overstate optimal marginal rates (Theorem 3 and Corollary 1), and the calibration in Section 12 shows that this overstatement is 1 to 4.5 percentage points across a wide range of plausible parameters. A government that designs its access policy ignoring the tax schedule will miss the feedback from optimal taxation to the income distribution that shapes the return to access investment.

13.2 What the Formulas Say About Tax Design in Practice

The over-taxation bias results have direct implications for empirical work on optimal tax rates. Much of the empirical literature on optimal taxation proceeds by estimating elasticities of taxable income with respect to

the net-of-tax rate and plugging them into the Saez [2001] formula [Chetty, 2009]. The present paper shows that this approach overstates optimal rates when $e_s > e_n$, that is, when income is increasing in the effective return $s = n \cdot m$ by more than it responds to the statutory rate n alone. The calibration provides the key message: for a moderate friction gradient of $\gamma = 0.10$, the bias is 1.5 to 2 percentage points on the top rate. For $\gamma = 0.20$, applicable in environments with steep access gradients, the bias rises to 3.5 to 4.5 percentage points. These are differences that matter for policy discussions about optimal top rate levels.

13.3 Access Policy as a Redistributive Instrument

Proposition 5 establishes that the social return to access investment rises when the income distribution deteriorates. The result bears on the *distributional value* of access, not only on aggregate take-up rates. When households at the bottom of the distribution face steep frictions (m low at low c), each unit of improvement in π generates a large welfare gain precisely for the agents whom redistribution is meant to help. The fact that this social return exceeds κ in equilibrium when inequality is high means that access policy is not only complementary to taxation but is itself a form of redistribution, one that works through the quality and efficiency of institutional intermediation rather than through the statutory tax code.

Table 2 makes this concrete. At the bottom of the income distribution, where the U-shaped rate schedule calls for high marginal rates to phase out transfers, friction corrections reduce optimal rates by roughly 23 per cent relative to the frictionless benchmark. This compression of the U-shape is not a paradox: it reflects the fact that high statutory rates at the bottom are partly undone by frictions, so the effective redistribution achieved at low income levels is already lower than the statutory schedule suggests. Improving access and reducing marginal rates at the bottom are complementary policy adjustments.

The participation tax result (Proposition 2 and Corollary 2) adds another layer to this picture. Programs like the EITC are designed on the premise that labor-market participation is elastically responsive to the financial gain from work, and the subsidy is sized accordingly. If access frictions reduce the effective participation response (because even employed workers face frictions in accessing the formal systems through which they receive the credit), then the EITC subsidy implicit in a friction-corrected optimal formula is smaller than standard analysis suggests. Conversely, if non-employment is heavily burdened by welfare program complexity, the gain from working is amplified relative to its statutory value, and the optimal subsidy is larger. The sign of the bias in EITC-style analysis is therefore an empirical question that depends on the relative magnitudes of frictions in employment and non-employment.

13.4 Limitations and Extensions

Several features of the model are deliberately simplified and could be generalized in future work. First, the access technology $\pi(a, c)$ is a reduced-form object. A more structural approach would model the specific frictions (time costs, information costs, stigma) and their dependence on household characteristics, allowing the optimal access policy to be differentiated across types. Second, the model treats a as a single aggregate instrument, and in practice governments choose among many components of administrative design. An extension with a vector of access instruments would allow for optimal targeting of different components of the administrative bundle. Third, the analysis is static. A dynamic version would allow access frictions to be accumulated as experience goods, so that households who navigate programs successfully in one period face lower frictions in subsequent periods. Fourth, the complementarity result in Proposition 4 is partial-equilibrium. A full general-equilibrium analysis would account for the effect of access policy on equilibrium wages and the labor supply of all agents, not just those directly affected by the access improvement. Fifth,

the calibration fixes γ from reduced-form evidence on take-up and administrative burdens rather than from a structural estimation. Future work estimating γ directly from individual-level data on income and effective consumption access would sharpen the quantitative conclusions.

14 Conclusion

The standard Mirrleesian optimal tax model assumes that disposable income translates directly and frictionlessly into welfare. This paper relaxes that assumption by introducing an access friction (a smooth wedge governed by the intensity function $\pi(a, c)$) between disposable income and effective consumption, and derives the consequences for optimal income tax design.

The main theoretical message is that access frictions modify the optimal tax problem in two simultaneous and reinforcing ways. On the welfare side, the social value of redistributing a dollar to a low-income household is lower when frictions are severe, because the dollar does not fully reach the household as consumption. On the incentive side, the effective return to earning is governed by the composite shifter $s = n \cdot m$ rather than the statutory net-of-tax rate alone, and when frictions rise with income (m increasing in c), the composite elasticity e_s exceeds the partial elasticity e_n. Both effects push the optimal marginal rate below the frictionless benchmark. Computations that ignore access frictions therefore overstate optimal rates at every income level throughout the distribution.

The calibration exercise translates this theoretical finding into quantitative terms. Using experimental evidence from Finkelstein and Notowidigdo [2019] on SNAP application costs, the income gradient of administrative burden documented in Herd and Moynihan [2018], and cross-program take-up data from Ko and Moffitt [2022], we calibrate the friction gradient γ at 0.05 for top earners, 0.10 for the broad working population, and 0.20 for environments with steep access gradients. Paired with standard elasticity and Pareto parameter values from Diamond and Saez [2011] and Saez et al. [2012], the resulting friction corrections are 1 to 4.5 percentage points on the optimal top rate. At the bottom of the income distribution, where frictions are most income-graded, the friction correction reduces optimal rates to about 77 per cent of the frictionless benchmark, compressing the U-shaped rate schedule from below.

Beyond the top bracket, the full-distribution formula (Theorem 2) shows that the over-taxation bias extends to the entire rate schedule under quasi-linear preferences. The participation tax formula (Proposition 2) shows that frictions introduce a further complication at the extensive margin: the optimal EITC-style subsidy depends on whether frictions are more severe for employed or non-employed agents, a distinction that varies across programs and contexts. The joint optimization result (Theorem 4 and Proposition 5) establishes that more unequal societies should invest more in access infrastructure, so that tax progressivity and access quality are complements in the optimal redistribution design.

Together, these results point toward a richer conception of optimal redistribution: one in which the institutional quality of public delivery systems is not a background constraint but a policy variable, jointly optimized alongside the tax schedule. Empirically, this suggests that studies of optimal taxation in environments with incomplete take-up (which encompass most real-world tax and transfer systems) should account for the access-friction wedge in both welfare weight construction and elasticity measurement. A practical implication for empirical researchers is that the statutory elasticity of taxable income, estimated from administrative tax data, understates the composite elasticity e_s, and that plugging it directly into the Saez formula yields top-rate prescriptions that are systematically too high by a margin the calibration shows to be economically significant.

References

Anthony B. Atkinson, Thomas Piketty, and Emmanuel Saez. Top incomes in the long run of history. *Journal of Economic Literature*, 49(1):3–71, 2011. doi: 10.1257/jel.49.1.3.

Timothy Besley and Stephen Coate. Workfare versus welfare: Incentive arguments for work requirements in poverty-alleviation programs. *American Economic Review*, 82(1):249–261, 1992.

Raj Chetty. Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1(1):451–488, 2009. doi: 10.1146/annurev.economics.050708.142910.

Janet Currie. The take-up of social benefits. In Alan Auerbach, David Card, and John Quigley, editors, *Public Policy and the Income Distribution*, pages 80–148. Russell Sage Foundation, 2006.

Peter Diamond and Emmanuel Saez. The case for a progressive tax: From basic research to policy recommendations. *Journal of Economic Perspectives*, 25(4):165–190, 2011. doi: 10.1257/jep.25.4.165.

Peter A. Diamond. Optimal income taxation: An example with a U-shaped pattern of optimal marginal tax rates. *American Economic Review*, 88(1):83–95, 1998.

Amy Finkelstein and Matthew J. Notowidigdo. Take-up and targeting: Experimental evidence from SNAP. *Quarterly Journal of Economics*, 134(3):1505–1556, 2019. doi: 10.1093/qje/qjz013.

Jon Gruber and Emmanuel Saez. The elasticity of taxable income: Evidence and implications. *Journal of Public Economics*, 84(1):1–32, 2002. doi: 10.1016/S0047-2727(01)00085-8.

Pamela Herd and Donald P. Moynihan. *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation, 2018.

Virginia Hernanz, Franck Malherbet, and Michele Pellizzari. Take-up of welfare benefits in OECD countries: A review of the evidence. OECD Social, Employment and Migration Working Papers 17, OECD Publishing, 2004.

Henrik Jacobsen Kleven and Wojciech Kopczuk. Transfer program complexity and the take-up of social benefits. *American Economic Journal: Economic Policy*, 3(1):54–90, 2011. doi: 10.1257/pol.3.1.54.

Wonsik Ko and Robert A. Moffitt. Take-up of social benefits. NBER Working Paper 30148, National Bureau of Economic Research, 2022. Revised March 2024.

James A. Mirrlees. An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208, 1971. doi: 10.2307/2296779.

Robert Moffitt. An economic model of welfare stigma. *American Economic Review*, 73(5):1023–1035, 1983.

Albert L. Nichols and Richard J. Zeckhauser. Targeting transfers through restrictions on recipients. *American Economic Review*, 72(2):372–377, 1982.

Thomas Piketty, Emmanuel Saez, and Stefanie Stantcheva. Optimal taxation of top labor incomes: A tale of three elasticities. *American Economic Journal: Economic Policy*, 6(1):230–271, 2014. doi: 10.1257/pol.6.1.230.

Emmanuel Saez. Using elasticities to derive optimal income tax rates. *The Review of Economic Studies*, 68 (1):205–229, 2001. doi: 10.1111/1467-937X.00166.

Emmanuel Saez. Optimal income transfer programs: Intensive versus extensive labor supply responses. *The Quarterly Journal of Economics*, 117(3):1039–1073, 2002. doi: 10.1162/003355302760193959.

Emmanuel Saez, Joel Slemrod, and Seth H. Giertz. The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of Economic Literature*, 50(1):3–50, 2012. doi: 10.1257/jel.50.1.3.

Nicolas Werquin. Income taxation with frictional labor supply. Manuscript dated April 2016 (Toulouse School of Economics), 2016.