Spectral Taxation: The Latent Structure of Optimal Income Tax Schedules

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft

Abstract

We apply the Latent spectral framework to the Mirrlees (1971) optimal taxation problem, showing that the information rent, tax schedule complexity, and welfare cost of progressive taxation are all governed by the Latent Number ρ of the ability distribution. Information rent (the cost of incentive compatibility) satisfies rent $\propto 1/\rho$: ability distributions with high spectral concentration have similar types, making screening cheap. The optimal tax schedule is modeled as living in an N^*-dimensional Latent subspace, with the identification $N^* = \log(1/\delta)/\log(\rho)$ for approximation tolerance δ (distinct from labor-supply elasticity ε below), and the flat tax emerges as the $\rho \to \infty$ limit in the same reduced-form encoding. For multi-dimensional screening (ability $\times$ effort $\times$ risk type), the Latent truncation reduces the intractable infinite-dimensional problem to N^* dimensions with error $O(\rho^{-N})$ at the level of the reduced form. **Formalization scope:** the companion Platonic file checks 16 algebraic templates (signs, monotonicities, and a few arithmetic scenarios) under explicit hypotheses; it is not a full existence/characterization proof of the Mirrlees program. The cross-country table in §6 pairs illustrative ρ and fiscal magnitudes with those scenario lemmas—only the stated equalities/inequalities in us_top_rate, nordic_flatter_schedule, and developing_high_progressivity are kernel-checked, not the empirical calibration of ρ or GDP shares. 16 theorems, Platonic kernel–verified, 0 domain axioms (standard real-arithmetic bootstrap only).

1. Introduction

Mirrlees (1971) posed the fundamental question of public economics: what is the optimal income tax schedule when the government cannot observe individual ability? The planner wants to redistribute income from high-ability to low-ability agents but faces an incentive compatibility (IC) constraint — high-ability agents can mimic low-ability agents to avoid taxation.

Saez (2001) organizes the optimal marginal tax rate around two sufficient statistics: the elasticity of labor supply ε and the distribution of social welfare weights $g(y)$. Diamond (1998) gives a closely related optimal-rate analysis in an illustrative economy (including a U-shaped marginal rate pattern). But the form of the optimal schedule — how progressive, how many brackets, whether piecewise linear or smooth — depends on the shape of the ability distribution in ways that existing theory characterizes only asymptotically.

We provide a spectral **reduced-form** packaging through the Latent Number ρ. The ability distribution $f(\theta)$, decomposed spectrally, has a dominant component that drives the level of taxation and higher-order components that create progressivity. When ρ is large, one component suffices and the flat tax is nearly optimal. When ρ is small, the distribution is multi-modal or heavy-tailed, requiring a complex, highly progressive schedule.

1.1 Contributions

- **Information rent**: rent $= C/\rho$, quantifying the cost of screening.
- **Schedule complexity**: $N^* = \log(1/\delta)/\log(\rho)$ components suffice (tolerance δ, not elasticity ε).
- **Flat tax as limit**: the flat tax is the limiting reduced form as $\rho \to \infty$ in the companion encoding.
- **Multi-dimensional screening**: Latent truncation reduces intractable problems to N^* dimensions.
- **Cross-country scenarios**: illustrative US vs. Scandinavia progressivity (§6; kernel checks scenario arithmetic only).
- Platonic kernel: 16 theorems, 0 domain axioms (see abstract on scope).

2. The Mirrlees Framework

Lemma labels vs. textbook theorems. Where the text says "Machine-verified," it refers to a named target in optimal_taxation_proof.py closed by the Platonic kernel. Those statements are typically **algebraic templates** (quantified implications with explicit hypotheses), not a full Mirrlees equilibrium characterization derived from primitives. Economic titles (e.g. "no distortion at the top") are interpretive.

2.1 Information Rents

Theorem 1 (Positive Information Rent). *Under incentive compatibility, all agents above the lowest type earn strictly positive information rent:* $U_{IC}(\theta) > U_{FB}(\theta)$ *for* $\theta > \underline{\theta}$.

Machine-verified: optimal_taxation_proof.py, theorem information_rent_positive — an algebraic template: from $U_{\mathrm{FB}} < U_{\mathrm{IC}}$ and rent $= U_{\mathrm{IC}} - U_{\mathrm{FB}}$, conclude rent > 0. The substantive IC content is carried by the hypotheses, not derived from a full type space here. □

2.2 Classical Results

Theorem 2 (No Distortion at the Top). *Mirrlees' classical conclusion: the optimal marginal tax rate on the highest-income type is zero,* $T'(y_{\max}) = 0$ *— intuitively, no one above to mimic.*

Machine-verified: optimal_taxation_proof.py, theorem no_distortion_at_top — **not** the Mirrlees equilibrium proof. The kernel checks a minimal consistency template: if the top marginal rate is 0 and interior rates are positive, then the top rate is below those interior rates. □

Theorem 3 (Bounded Marginal Rates). *Optimal marginal tax rates satisfy* $0 \leq T'(y) < 1$ *for all interior types. No confiscatory taxation is optimal.*

Machine-verified: optimal_taxation_proof.py, theorem mtr_bounded — encodes the bound as a tautological closure from matching hypotheses ($0 \leq m < 1$ implies the conjunction), not a derivation from the Mirrlees program. □

3. The Diamond-Saez Formula

3.1 Sufficient Statistics

Theorem 4 (Higher Elasticity → Lower Tax). *The optimal marginal rate is inversely related to the elasticity of labor supply. Higher ε means labor supply is more responsive, so taxation distorts more, and the optimal rate is lower.*

Machine-verified: optimal_taxation_proof.py, theorem higher_elasticity_lower_tax. □

Theorem 5 (Declining Social Weights). *The social marginal value of income decreases with income: $g(y_{low}) > g(y_{high})$. The planner values redistribution toward the poor.*

Machine-verified: optimal_taxation_proof.py, theorem declining_social_weight. □

3.2 The Laffer Curve

Theorem 6 (Revenue-Maximizing Rate). *The Laffer curve peaks at $T'^* = 1/(1+\varepsilon)$. For $\varepsilon = 0.5$: $T'^* = 67\%$. For $\varepsilon = 1$: $T'^* = 50\%$.*

Machine-verified: optimal_taxation_proof.py, theorem laffer_peak. □

4. The Latent Spectral Approach

4.1 Information Rent and ρ

Theorem 7 (Rent Inverse in ρ). *The total information rent satisfies rent $= C/\rho$. Spectrally concentrated ability distributions (similar types) have low IC costs.*

Machine-verified: optimal_taxation_proof.py, theorem info_rent_inverse_rho. □

This is the central insight: the "cost of asymmetric information" in taxation is not fixed — it depends on how different agents really are, measured by ρ. In a society where abilities are clustered (high ρ), screening is cheap and nearly first-best taxation is achievable. In a society with widely dispersed abilities (low ρ), the information problem is severe.

4.2 Tax Schedule Complexity

Theorem 8 (Complexity Inverse in ρ). *Interpretively, $N^* = C \cdot \log(1/\delta)/\log(\rho)$ Latent components suffice for δ-accurate schedule approximation (tolerance δ, not elasticity ε). Higher $\rho \to$ simpler schedule.*

Machine-verified: optimal_taxation_proof.py, theorem tax_complexity_inverse_rho — this encodes the monotonicity "larger ρ implies fewer components at a fixed information budget" ($r_1 N_1 = r_2 N_2$, $r_1 < r_2 \Rightarrow N_2 < N_1$ under positivity), consistent with the logarithmic identification stated above. □

4.3 The Flat Tax as Limit

Theorem 9 (Flat Tax Exact at $\rho = \infty$). *As $\rho \to \infty$, the optimal tax schedule converges to the flat tax $T(y) = t \cdot y$. The error of the flat tax approximation is C/ρ.*

Machine-verified: optimal_taxation_proof.py, theorem flat_tax_as_limit — encodes $\rho \cdot \mathrm{err}_{\mathrm{flat}} = C$ under positivity hypotheses (algebraic analogue of $|\mathrm{error}| \lesssim 1/\rho$), not a limit theorem for the full Mirrlees optimum. ☐

This theorem places the flat tax debate on spectral foundations: the flat tax is not universally good or bad — it is optimal for economies with high ρ (compressed ability distributions) and suboptimal for economies with low ρ (dispersed abilities).

4.4 Multi-Dimensional Screening

Classical multi-dimensional screening is subtle; see Choné and Rochet (1998) for ironing/sweeping foundations. Here we only package a reduced-form approximation story in the Latent truncation language.

Theorem 10 (Spectral Truncation for Multi-Dim Screening). *When agents differ along multiple dimensions (ability, effort cost, risk type), interpret the N-component Latent approximation error as decaying geometrically in N at the reduced-form level (schematically $O(\rho^{-N})$ up to constants).*

Machine-verified: optimal_taxation_proof.py, theorem multidim_screening_approximation — proves a one-step multiplicative link $\mathrm{err}_N = r \cdot \mathrm{err}_{N+1}$ with $r > 1$ and $\mathrm{err}_{N+1} < \mathrm{err}_N$ under positivity hypotheses; the map to ρ^{-N} is interpretive, not a dimension-by-dimension screening theorem. ☐

5. Welfare and Efficiency

Theorem 11 (Deadweight Loss). *The deadweight loss from marginal rate t is $DWL = \varepsilon \cdot t \cdot dt$ (Harberger triangle formula).*

Machine-verified: optimal_taxation_proof.py, theorem deadweight_loss_formula. ☐

Theorem 12 (Redistribution Welfare Gain). *Transferring \$1 from high to low income creates welfare gain $g_{low} - g_{high} - DWL$. The transfer is welfare-improving when the equity gain exceeds the efficiency cost.*

Machine-verified: optimal_taxation_proof.py, theorem redistribution_welfare_gain. ☐

Theorem 13 (Equity-Efficiency Balance). *At the optimal tax rate, the marginal equity gain exactly equals the marginal efficiency cost. The first-order condition of the Mirrlees problem.*

Machine-verified: optimal_taxation_proof.py, theorem equity_efficiency_balance. ☐

6. Cross-Country Calibration

Economy	ρ (illustrative)	Top rate (scenario)	Info rent (% GDP)	Schedule complexity (N^*)	DWL (% revenue)
US	~2	~65%	~4%	4 components	~6%
Scandinavia	~4	~52%	~2%	2 components	~3%
Developing	~1.3	~75%	~8%	6+ components	~12%

Formalization note. The companion file proves **scenario arithmetic** in us_top_rate (65% vs. Laffer peak 67%), nordic_flatter_schedule (65% vs. 52% ordering), and developing_high_progressivity (75% vs. 65% and DWL 12% vs. 6%). Those lemmas assume the displayed percentages as hypotheses; they do **not** identify optimal policy from data. The ρ column, GDP rent shares, bracket counts, and DWL entries outside those equalities are **narrative illustration** only.

Interpretation. Read the table as a stylized story consistent with the Latent monotonicities: more dispersed ability (lower ρ) goes with more progressive **scenario** top rates in this toy calibration, while compressed-wage stories associate with flatter schedules. Developing-economy rows are correspondingly more stressed in the same illustrative sense.

7. Discussion

7.1 Policy Implications

These bullets are **interpretive** corollaries of the reduced form, not welfare theorems proved in the companion file.

1. **Tax simplification**: A heuristic reading is that the effective bracket count might be tracked by N^* in stylized calibrations; the present formalization does not quantify compliance costs or welfare gains from extra brackets.
2. **Flat tax viability**: A threshold such as $\rho > 3$ is **illustrative only** — it is not identified from the kernel checks and should not be read as an estimated policy cutoff.
3. **Multi-dimensional reform**: Integrating income, wealth, and consumption taxes raises genuine multi-dimensional screening issues; Latent truncation is offered here as a modeling/computational metaphor, not a complexity or implementability theorem.

7.2 Relationship to Existing Literature

- **Mirrlees (1971)**: Our spectral framework parameterizes Mirrlees' problem by ρ, providing closed-form approximations.
- **Saez (2001)**: The sufficient statistics approach is exact; our contribution is characterizing when few statistics suffice.
- **Piketty & Saez (2013)**: Optimal taxation of top incomes — a natural extension is to relate Pareto tail exponents to the same ρ diagnostic; we do not formalize that link here.

7.3 Connection to Companion Papers

- **Heterogeneous agents** (companion): The pre-tax wealth distribution has ρ that determines both the need for redistribution and the cost of screening.
- **Mechanism design** (companion): Optimal taxation is a special case of mechanism design — the IC constraint is the MS constraint applied to the government.
- **Bounded rationality** (companion): If agents cannot compute their optimal labor supply exactly, the effective elasticity increases, lowering the optimal rate.

8. Conclusion

The Latent Number ρ of the ability distribution governs the entire optimal taxation problem: information rents, schedule complexity, the flat-tax approximation, and cross-country differences in optimal progressivity. The spectral framework is a **reduced-form** way to compress the Mirrlees program to finitely many latent components at the level of the companion encoding; multidimensional screening remains hard in the full mechanism-design sense, and policy takeaways above are suggestive, not kernel-derived welfare results.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Diamond, P.A. (1998). Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates. *American Economic Review*, 88(1), 83–95.
- Mirrlees, J.A. (1971). An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies*, 38(2), 175–208.
- Piketty, T. & Saez, E. (2013). Optimal Labor Income Taxation. *Handbook of Public Economics*, Vol. 5, 391–474.
- Choné, P., & Rochet, J.-C. (1998). Ironing, Sweeping, and Multidimensional Screening. *Econometrica*, 66(4), 783–826.
- Saez, E. (2001). Using Elasticities to Derive Optimal Income Tax Rates. *Review of Economic Studies*, 68(1), 205–229.