Latent Mechanism Design: Spectral Approximation of Optimal Mechanisms

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft

Abstract

We apply the Latent spectral framework to mechanism design, framing bilateral-trade tensions through a spectral parameter ρ (Latent Number) of the type distribution. The VCG mechanism achieves full efficiency but runs a budget deficit; Myerson and Satterthwaite (1983) proved that no mechanism simultaneously achieves efficiency, incentive compatibility, individual rationality, and budget balance. Machine-checked lemmas include stepwise geometric decay of a formal approximation error under a multiplicative recurrence (Theorem 5), explicit $\rho = 2$ error normalizations (Theorem 6), and a gap comparison across coupled (r, Δ) scalings (Theorem 7). For $\rho = 2$, Theorem 6 gives five-component approximation error $\varepsilon_5 = C/32$; the verified numerics use spectral error $\varepsilon = 1/8$ and the uniform bilateral-trade gap calibration $1/4$, yielding efficiency $\geq 7/8$ (87.5%) and a deficit bound $1/32 \approx 3.1\%$ of full surplus scale, with strictly lower deficit than a VCG benchmark under the hypotheses of latent_beats_vcg_deficit. 16 theorems, machine-verified in Platonic (0 user axioms in the companion file).

1. Introduction

Mechanism design — the engineering of rules for strategic interaction — faces a fundamental impossibility: Myerson and Satterthwaite (1983) showed that for bilateral trade with private valuations, no mechanism can simultaneously be:

1. **Efficient** (maximize total surplus)
2. **Incentive compatible** (truthful reporting is optimal)
3. **Individually rational** (voluntary participation)
4. **Budget balanced** (no external subsidy)

The VCG mechanism achieves (1)–(3) but violates (4), running a deficit proportional to the "information rent" agents extract. The question is: how close can we get to all four properties simultaneously?

We organize the tradeoff through a spectral viewpoint: the Latent Number ρ of the joint type distribution (and related truncation-error parameters in the formal layer). **The machine-checked fragment below is a scalar surrogate**—real arithmetic lemmas under explicit hypotheses—so narrative claims about ρ should be read as aligned with those hypotheses, not as a self-contained game-theoretic characterization of all mechanisms. Heuristically, sharper spectral concentration is associated with smaller formal gaps in the coupled (r, Δ) calibration (Theorem 7 and §3.3); the uniform-type calibration in §5.1 illustrates a **large** formal gap at the 1/4 normalization.

1.1 Contributions

- **Latent spectral approximation**: Theorem 5 gives strict stepwise decay of formal errors under a ratio recursion; Theorem 6 fixes $\rho = 2$ normalizations $\varepsilon_3 = C/8$, $\varepsilon_5 = C/32$.
- **MS / gap comparisons**: Theorem 3 gives positivity under overlap hypotheses; Theorem 7 compares gaps across r-scalings at fixed product $r\Delta$.
- **Illustrative near-optimality numerics**: for the $\rho = 2$ normalization in Theorem 6 and the spectral error $\varepsilon = 1/8$, the companion file derives efficiency $\geq 7/8$ and deficit bound $1/32$ **once** the $1/4$ gap calibration is adjoined (latent_efficiency_87pct, latent_deficit_1_32); five components satisfy $\varepsilon_5 = C/32$ (Theorem 6).
- **VCG comparison**: under the formal comparison hypotheses, the Latent construction has strictly smaller deficit than a VCG-style benchmark (latent_beats_vcg_deficit).
- Machine-verified: 16 theorems in mechanism_design_proof.py, 0 domain p.axiom declarations (bootstrap + tactics only).

2. Preliminaries

2.1 The VCG Mechanism

In the VCG (Vickrey-Clarke-Groves) mechanism, each agent pays the externality they impose on others. This ensures:

Theorem 1 (VCG Efficiency — formal encoding). *In the surplus-ordering model checked in vcg_efficiency, the optimal surplus weakly dominates any alternative allocation surplus.* Machine-verified. □

Theorem 2 (VCG Budget Deficit — sufficient condition). *Under the hypotheses of vcg_budget_deficit (cost exceeds total payments, with deficit defined accordingly), the budget imbalance is strictly positive.* Machine-verified. □

2.2 The Myerson-Satterthwaite Impossibility

Theorem 3 (MS Gap — positivity in the formal model). *Under the overlap hypotheses encoded in ms_gap_positive (buyer–seller support overlap with a positive formal gap bounded by overlap), the gap is strictly positive.*

Machine-verified: mechanism_design_proof.py, theorem ms_gap_positive. □

Theorem 4 (Efficiency–deficit tension). *Fix a positive impossibility gap and a deficit tolerance strictly between* 0 *and the gap. If efficiency loss is modeled as gap minus deficit tolerance, then that loss is strictly positive and strictly smaller than the gap.*

Machine-verified: mechanism_design_proof.py, theorem efficiency_budget_tradeoff. □

3. The Latent Spectral Approach

3.1 Type Distribution Decomposition

The joint type distribution $F(v_1, v_2)$ admits a Latent spectral decomposition with eigenvalues $\sigma_1 \geq \sigma_2 \geq \cdots$ and Latent Number $\rho = \sigma_1/\sigma_2$.

3.2 Truncation Approximation

Theorem 5 (Spectral Approximation Decay). *If successive approximation errors satisfy* $r\,\varepsilon_{n+1} = \varepsilon_n$ *with* $r > 1$ *and* $\varepsilon_n > 0$*, then* $\varepsilon_{n+1} < \varepsilon_n$ *(strict geometric decrease stepwise).*

Machine-verified: mechanism_design_proof.py, theorem spectral_approximation_decay. □

Theorem 6 (Rapid Convergence for $\rho = 2$). *If* $8\varepsilon_3 = C$ *and* $32\varepsilon_5 = C$ *with* $C > 0$*, then* $\varepsilon_5 < \varepsilon_3$ *(five components give a strictly smaller error than three at this normalization).*

Machine-verified: mechanism_design_proof.py, theorem rapid_convergence_rho2. □

3.3 The MS Gap in Latent Terms

Theorem 7 (Gap comparison across ρ-scalings). *If* $r_1\Delta_1 = r_2\Delta_2$ *with* $0 < r_1 < r_2$ *and both gaps positive, then* $\Delta_2 < \Delta_1$ *(the gap attached to the larger* r*-scale is smaller).*

Machine-verified: mechanism_design_proof.py, theorem latent_gap_inverse_rho. □

Narrative link: Interpreting r as a spectral-concentration proxy, this formal lemma supports the story that sharper concentration is associated with smaller formal gaps in the coupled calibration; the quantitative $\Delta \propto 1/\rho$ headline is motivational, not a literal output of the current file.

3.4 Finite-Dimensional Reduction

Theorem 8 (Finite-dimensional feasibility slice). *For any* $0 < \varepsilon < C$ *and any* $r > 1$*, one has* $C - \varepsilon > 0$ *(a positivity condition used when truncating at finite* ε*).*

Machine-verified: mechanism_design_proof.py, theorem finite_dimensional_reduction. □

4. The Latent Near-Optimal Mechanism

4.1 Properties

Theorem 9 (Near-Optimality — multiplicative form). *Under the hypotheses of latent_mechanism_near_optimal, latent efficiency equals* $(1 - \varepsilon)$ *times a positive baseline optimal efficiency, with* $0 < \varepsilon < 1$*.*

Theorem 10 (Bounded Deficit — product form). *Under latent_deficit_bounded, deficit is modeled as gap times an interior spectral error parameter* $\varepsilon \in (0, 1)$*, yielding a strictly positive deficit strictly below the gap.*

Both machine-verified in mechanism_design_proof.py. □

4.2 Dominance over VCG

Theorem 11 (Latent Beats VCG on Deficit). *Under the hypotheses of latent_beats_vcg_deficit (including a nonnegative VCG-style deficit at least as large as the impossibility gap and an interior spectral error* ε*), the Latent deficit is strictly smaller.*

Machine-verified: mechanism_design_proof.py, theorem latent_beats_vcg_deficit. □

The intuition: VCG ignores the spectral structure of types, paying full information rents to all agents. The Latent mechanism exploits the concentration of types in spectral space, extracting rents only along the dominant components.

5. Numerical Examples

5.1 Uniform bilateral trade (unconcentrated calibration)

In the calibration used for uniform_ms_gap, the formal surrogate gap satisfies $4\,\Delta_{\text{MS}} = 1$, i.e. $\Delta_{\text{MS}} = 1/4$ (25% at this normalization). We use this as the **diffuse / unconcentrated** benchmark in the storytelling; we do **not** compute a literal eigenvalue ratio ρ for $U[0,1]^2$ in the companion file. Interpreting Δ_{MS} as surplus-scale tension at this normalization, the obstruction is large.

Machine-verified: mechanism_design_proof.py, theorem uniform_ms_gap. □

5.2 Concentrated Types ($\rho = 2$): verified numerics

Theorems latent_efficiency_87pct and latent_deficit_1_32 use spectral error $\varepsilon = 1/8$ (so efficiency $\eta = 1 - \varepsilon = 7/8 = 87.5\%$) together with the same 1/4 gap calibration, giving deficit bound $1/32 \approx 3.125\%$ at that scale. Separately, Theorem 6 (rapid_convergence_rho2) compares $\varepsilon_3 = C/8$ and $\varepsilon_5 = C/32$ for five versus three components when $C > 0$.

Machine-verified: mechanism_design_proof.py, theorems latent_efficiency_87pct, latent_deficit_1_32, rapid_convergence_rho2. □

6. Discussion

6.1 Applications

The Latent mechanism design framework applies to: - **Spectrum auctions**: bidder valuations have low-dimensional spectral structure (geographic/demographic factors). - **Matching markets**: preferences decompose into a few dominant factors. - **Carbon trading**: emission valuations concentrate around industry-specific benchmarks.

In each case, the designer can estimate ρ from historical data and choose N to balance efficiency against budget balance.

6.2 Relationship to Existing Work

- **Myerson (1981)**: optimal auction theory in private-value models. This note does **not** extend that framework to multi-dimensional types; the Platonic lemmas manipulate **scalar** surrogates (surplus levels, payments, gaps).
- **Crémer & McLean (1988)**: full surplus extraction under correlation. Our ρ is a **Latent spectral / truncation parameter in the formal encoding**, not a redefinition of their correlation conditions.
- **Bergemann & Välimäki (2010)**: dynamic pivot mechanisms. A genuine multi-period Latent truncation story would be **conjectural here**—not part of mechanism_design_proof.py.

6.3 Limitations

- The spectral decomposition assumes the designer knows (or can estimate) the type distribution.
- Computational tractability requires N to be moderate; for very low ρ, many components may be needed.

- **Formalization scope:** Machine verification is in Platonic real arithmetic; several theorems are algebraic encodings of economic claims (explicit hypotheses in mechanism_design_proof.py). The narrative should be read as aligned with those types, not as a substitute for a full game-theoretic manuscript.
- **What this paper does not claim:** A complete independent proof of the Myerson–Satterthwaite theorem, closed-form optimal auctions beyond the cited literature, or empirical estimates of ρ from data.

7. Conclusion

The Latent Number ρ and coupled scalings (r, Δ) organize **formal** surrogates for efficiency–deficit tension in the companion proof file. Larger spectral-concentration proxies align with **smaller** formal gaps in the coupled calibration (Theorem 7); they do **not**, by themselves, establish existence of exact ex post budget-balanced mechanisms at full efficiency. The Latent truncation narrative is a **constructive lens** for finite-dimensional approximation; the verified core is the inequality layer in mechanism_design_proof.py.

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

- Bergemann, D. & Välimäki, J. (2010). The Dynamic Pivot Mechanism. *Econometrica*, 78(2), 771–789.
- Clarke, E.H. (1971). Multipart Pricing of Public Goods. *Public Choice*, 11(1), 17–33.
- Crémer, J. & McLean, R.P. (1988). Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions. *Econometrica*, 56(6), 1247–1257.
- Groves, T. (1973). Incentives in Teams. *Econometrica*, 41(4), 617–631.
- Myerson, R.B. (1981). Optimal Auction Design. *Mathematics of Operations Research*, 6(1), 58–73.
- Myerson, R.B. & Satterthwaite, M.A. (1983). Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory*, 29(2), 265–281.
- Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1), 8–37.