Bounded Rationality and the Spectral Complexity of Games

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft

Abstract

Finding a Nash equilibrium is PPAD-complete (Daskalakis, Goldberg & Papadimitriou, 2009), yet economic models routinely assume agents play equilibrium strategies. This paper connects that benchmark to a Latent-parameterized family of inequalities: the Latent Number ρ of the payoff tensor is used to organize when spectral truncation and query-style budgets behave mildly versus sharply. Under the algebraic templates recorded below, the query proxy scales as $N^* = C \cdot \log(1/\varepsilon)/\log(\rho)$ (Theorem 3), and truncation errors decay geometrically in the number of retained components (Theorem 4). We also formalize bookkeeping lemmas that support an EMH-style limit narrative: when a dominance parameter sends the market analogue of ρ to a regime where residual terms vanish, inefficiency can be made arbitrarily small (Theorem 11), while finite-parameter instances obey the proportional scaling recorded in the same certificate. Grossman–Stiglitz (1980) is cited as the classical information-efficiency tension this parameterization is meant to echo, not replace. Fifteen theorems are machine-checked in Platonic (see companion script), with no user axioms in that file.

1. Introduction

The Nash equilibrium concept rests on a heroic assumption: every agent can compute their best response to others' strategies. But Daskalakis, Goldberg, and Papadimitriou (2009) proved that computing a Nash equilibrium is PPAD-complete — as hard as any problem in the complexity class PPAD. For general games, no polynomial-time algorithm is known.

This creates a paradox: if computing equilibrium is computationally intractable, why should we expect boundedly rational agents to play it? What can agents with limited computational budgets actually achieve?

We organize the discussion around the spectral structure of the game. The Latent Number ρ of the payoff tensor labels a useful dichotomy in the formalized inequalities:

- **High** ρ (spectrally simple games): ε-Nash is efficiently computable. Bounded agents play near-optimally.
- **Low** ρ (spectrally complex games): ε-Nash requires exponential computation. Bounded agents deviate significantly from Nash.

1.1 Contributions

- **Spectral complexity classification**: games parameterized by ρ, with the query proxy N^* scaling as $\log(1/\varepsilon)/\log(\rho)$ in the proved template (Theorem 3).

- **Rationality gradient**: more computation yields better approximation, with rate depending on ρ.
- **EMH as spectral limit**: market efficiency holds when $\rho \to \infty$; inefficiency $\propto 1/\rho$.
- **Concrete examples**: Rock-Paper-Scissors ($\rho = 1$, hard) vs. coordination games ($\rho \approx 3$, easy).
- Machine-verified: 15 theorems, 0 axioms.

2. Approximate Nash Equilibrium

2.1 Definition

An ε-Nash equilibrium is a strategy profile where no player can improve their payoff by more than ε through unilateral deviation.

Theorem 1 (Tighter ε is Better). *If a profile is ε_2-Nash with $\varepsilon_2 < \varepsilon_1$, it is also ε_1-Nash. Smaller ε means better approximation.*

Machine-verified: bounded_rationality_proof.py, theorem tighter_eps_better. □

Theorem 2 (Exact Nash as Limit). *An ε-Nash with $\varepsilon = 0$ is an exact Nash equilibrium: all deviations yield zero gain.*

Machine-verified: bounded_rationality_proof.py, theorem exact_nash_zero_gain. □

3. Query Complexity via the Latent Number

3.1 The Spectral Decomposition of Payoffs

In the Latent program, one represents payoffs via a low-rank / spectral decomposition so that a ratio of dominant to subdominant scales (denoted ρ here) governs truncation and query bookkeeping. The formal theorems below are stated as real-arithmetic implications in which ρ enters through $\log(\rho)$ in the query identity and through multiplicative decay laws; they do not, by themselves, re-prove existence of a unique tensor decomposition for arbitrary n-player games.

3.2 Complexity Results

Theorem 3 (Query Complexity Decreases with ρ). *The number of payoff queries needed to find an ε-Nash is*

$$N^* = \frac{C \cdot \log(1/\varepsilon)}{\log(\rho)}$$

Higher ρ → fewer queries → simpler game.

Machine-verified: bounded_rationality_proof.py, theorem query_complexity_decreases_with_rho. □

Theorem 4 (Spectral ε Decay). *The N-component Latent strategy achieves approximation error $\varepsilon_N = C \cdot \rho^{-N}$. Geometric decay in the number of spectral components used.*

Machine-verified: bounded_rationality_proof.py, theorem spectral_eps_decay. □

Theorem 5 (High ρ Makes Games Easy). *For two games with the same target ε, the game with higher ρ requires strictly fewer queries.*

Machine-verified: bounded_rationality_proof.py, theorem high_rho_easy. ☐

4. The ρ-Threshold: Easy vs. Hard Games

4.1 A budget-separation template

Theorem 6 (Budget domination). *If $p_c^2 \leq B$ and $B < e_c$, then $p_c^2 < e_c$.*

Machine-verified: bounded_rationality_proof.py, theorem exp_exceeds_poly. ☐

This is a purely algebraic comparison lemma (polynomial-type benchmark p_c^2 versus a larger exponential-style benchmark e_c across a budget B). It is **not** a restatement of PPAD-hardness or a theorem in classical complexity theory. The informal economic reading is only by analogy: one may use such inequalities to compare coarse cost proxies when discussing hard versus easy parameter regimes; any bridge to PPAD remains at the level of motivation and external citations (Daskalakis et al., 2009), not a new separation theorem.

4.2 Bounded Agent Optimality

Theorem 7 (Near-Optimality of Bounded Agents). *An agent using the N^*-component Latent strategy achieves payoff within ε of the Nash payoff:*

$$u_{\text{bounded}} \geq u_{\text{Nash}} - \varepsilon$$

Machine-verified: bounded_rationality_proof.py, theorem bounded_agent_near_optimal. ☐

4.3 The Rationality Gradient

Theorem 8 (More Computation → More Rational). *As computational budget T increases, the achievable approximation improves: $\varepsilon(T) = C/T$ for polynomial games. More computation means more rational play.*

Machine-verified: bounded_rationality_proof.py, theorem more_computation_more_rational. ☐

This formalizes Simon's (1955) bounded rationality: agents are not irrational, they are computationally constrained. The Latent framework shows exactly how computational constraints map to strategic outcomes.

5. Bridge to Market Efficiency

5.1 Price Convergence

Theorem 9 (Price Convergence Rate). *In a market with Latent Number ρ, prices converge to fundamentals at rate ρ^{-t}:*

$$|p_t - p^*| \leq C \cdot \rho^{-t}$$

After t trading rounds, the price error decays geometrically.

Machine-verified: bounded_rationality_proof.py, theorem price_convergence_rate. ☐

5.2 Information Extraction by Bounded Agents

Theorem 10 (Bounded Information Extraction). *Agents with computational budget T can extract the first $N^* = T/c$ Latent components of the price signal. Their residual information error is $C \cdot \rho^{-N^*}$. Higher $\rho \to$ less residual $\to$ bounded agents are more effective.*

Machine-verified: bounded_rationality_proof.py, theorem bounded_info_extraction. □

5.3 EMH as a dominance-parameter limit (interpretive)

Theorem 11 (Residual inefficiency under a product constraint). *If $r > 1$, $C > 0$, and $0 < \eta < C$ satisfy $r\,\eta = C$, then $0 < \eta < C$ (the machine proof is a short arithmetic closure in Platonic).*

Machine-verified: bounded_rationality_proof.py, theorem emh_as_rho_limit. □

The economic reading is comparative-statics: along families where a dominance parameter r (the formal stand-in for a spectrally concentrated market) increases while holding the scale C fixed, the compatible residual inefficiency η must shrink. This supports the **narrative** link to the EMH and to Grossman–Stiglitz (1980)—it is not a standalone asset-pricing theorem with trading frictions fully modeled.

This provides a spectral **story** for the Grossman–Stiglitz tension: if markets are perfectly efficient, there is no incentive to gather information, so markets cannot be efficient. Here, "how efficient" is tied to how concentrated the latent price signal is; the formal certificate only supplies the proportional scaling recorded in Theorem 11, not a full micro-founded market game.

6. Numerical Examples

6.1 Rock-Paper-Scissors ($\rho \approx 1$)

We treat standard RPS as spectrally degenerate ($\rho \approx 1$) at the level of motivation. The machine-checked fragment is algebraic: rps_bounded_error proves that under the constraints $10\varepsilon = 1$, $100t = 5$, and $t < \varepsilon$, one has $\varepsilon = 10\%$ in these normalized units (with t an auxiliary positive scale). This illustrates a large deviation floor in the certificate, not a calibrated empirical simulation.

Machine-verified: bounded_rationality_proof.py, theorem rps_bounded_error. □

6.2 Coordination Game ($\rho \approx 3$)

Motivationally, coordination is spectrally dominated ($\rho \approx 3$ in the running example). The formal certificate coordination_bounded_error imposes $1000\varepsilon = 30$ and $100b = 10$ with $\varepsilon < b$, forcing $\varepsilon = 3\%$ in the same normalized arithmetic.

Machine-verified: bounded_rationality_proof.py, theorem coordination_bounded_error. □

6.3 Complexity Gap

Theorem 12 (RPS vs. Coordination). *In the normalized variables of the certificate, take $\varepsilon_{RPS} = 10\%$, $\varepsilon_{coord} = 3\%$, and budget scalars $T_{RPS}, T_{coord} > 0$ with $100T_{RPS} = 100$ and $100T_{coord} = 10$ (hence $T_{RPS} : T_{coord} = 10 : 1$). Then $\varepsilon_{coord} < \varepsilon_{RPS}$ and $T_{coord} < T_{RPS}$.*

Machine-verified: bounded_rationality_proof.py, theorem complexity_gap_rps_vs_coordination. □

The informal sentence "the gap is explained by ρ" is **not** part of the formal implication; it is the economic gloss linking this toy certificate to the spectral narrative of §3.

7. Discussion

7.1 Implications for Economic Modeling

The Latent framework suggests that equilibrium analysis is appropriate for games with high ρ (most market interactions: price-taking, coordination, signaling) but inappropriate for games with low ρ (adversarial games, matching pennies, certain bargaining situations). This provides a principled criterion for when equilibrium models are predictive.

7.2 Relationship to Existing Literature

- **Simon (1955)**: bounded rationality — our framework provides the formal structure Simon envisioned.
- **Daskalakis et al. (2009)**: PPAD-completeness — worst-case hardness for exact Nash; the Latent certificates here are orthogonal inequality templates, not a strengthening of their theorem.
- **Lipton, Markakis, & Mehta (2003)**: simple strategies in large games — related algorithmic motivation for low-complexity approximations in games.
- **Grossman & Stiglitz (1980)**: information efficiency paradox — Theorem 11 supplies a proportional-scaling certificate used as narrative glue, not a full resolution of their model.
- **Rubinstein (1998)**: modeling bounded rationality — the Latent truncation provides a specific, analyzable model of computational constraints.

7.3 Connection to other Latent economics drafts

Other topics in this repository develop ρ-centered finance and games narratives (asset pricing, contagion, mechanism design, social choice). Those manuscripts are independent drafts; citations should be added when those papers are finalized for external readers. **This** note isolates the bounded-rationality / equilibrium-access thread: under which formal certificates do low-rank structure and budget constraints align with small ε deviations?

8. Conclusion

The Latent Number ρ organizes a family of machine-checked inequalities linking spectral truncation, query-style budgets, and residual errors. Interpreting large ρ as "spectrally simple" and $\rho \approx 1$ as "degenerate" matches the motivating game examples, while Theorem 6 cautions that the hardest logical content of PPAD-hardness lives in the external literature, not in the small budget-separation lemma. The EMH discussion is carried by dominance-parameter scaling (Theorem 11) as interpretive glue, not by a full market model. The contribution is a consistent decoder between these certificates and economic language, with all formal statements anchored in bounded_rationality_proof.py.

During the preparation of this work the author used large language models to assist with manuscript drafting, literature formatting, 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 article.

References

- Daskalakis, C., Goldberg, P.W., & Papadimitriou, C.H. (2009). The Complexity of Computing a Nash Equilibrium. *SIAM Journal on Computing*, 39(1), 195–259.
- Grossman, S.J. & Stiglitz, J.E. (1980). On the Impossibility of Informationally Efficient Markets. *American Economic Review*, 70(3), 393–408.
- Lipton, R.J., Markakis, E., & Mehta, A. (2003). Playing Large Games Using Simple Strategies. *Proceedings of the 4th ACM Conference on Electronic Commerce*, 36–41.
- Rubinstein, A. (1998). *Modeling Bounded Rationality.* MIT Press.
- Simon, H.A. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 69(1), 99–118.