Quantitative Arrow: Measuring Distance from Impossibility via the Latent Framework

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft

Executive summary (non-technical)

Arrow's theorem is a *compatibility* statement: three familiar fairness axioms cannot all hold exactly once there are at least three alternatives. That leaves a practical question—how much slack a real aggregation rule needs—and a modeling question: whether "almost" satisfying the axioms is easier when preferences line up along a low-dimensional latent structure. This note adopts the Latent program's spectral concentration parameter ρ as a storytelling device: it sketches how modeled stress on independence-style conditions might scale like $1/\rho^2$ while dictator-adjacency templates scale like $1/\rho$, and it pairs that narrative with a dozen machine-checked *scalar* inequalities that keep the quantitative claims aligned with what the proof kernel actually certifies. The classical impossibility for $m \geq 3$ is not overturned; the contribution is honest bookkeeping between folklore economics, interpretive heuristics, and verified real arithmetic.

Abstract

Arrow's Impossibility Theorem (1951) proves that no social welfare function over three or more alternatives can simultaneously satisfy Unanimity, Independence of Irrelevant Alternatives (IIA), and Non-Dictatorship. This paper develops a **quantitative narrative** using the Latent Number ρ: (1) IIA violation intensity is modeled to scale like $1/\rho^2$, (2) dictator-adjacency is modeled to scale like $1/\rho$, and (3) a companion scalar encoding records a manipulation gain bound compatible with $1/(N-1)$ when a linear constraint $(N-1)g \leq 1$ is imposed. The Latent viewpoint suggests a continuous "distance from Arrow" reading of aggregation stress, complementary to the classical impossibility statement. For an illustrative normalization with $C = 1$ and $\rho = 3$, the toy rate $C/\rho^2 = 1/9$ is $\approx 11\%$. Twelve lemmas in arrow_impossibility_proof.py are machine-checked in the Platonic kernel (real-arithmetic scaffolding; no extra domain axioms beyond the standard bootstrap).

1. Introduction

Arrow's Impossibility Theorem is one of the most celebrated results in mathematical social science. It states that for three or more alternatives, no social welfare function $f : \mathcal{L}^N \to \mathcal{L}$ (mapping profiles of linear orders to a social ranking) can satisfy:

1. **Unanimity (U)**: If all voters rank $a > b$, then society ranks $a > b$.
2. **Independence of Irrelevant Alternatives (IIA)**: The social ranking of a vs. b depends only on individual rankings of a vs. b.
3. **Non-Dictatorship (ND)**: No single voter determines the social ranking.

The theorem has been interpreted as showing the impossibility of "fair" aggregation of preferences. But this interpretation is too strong: the theorem says these three axioms are jointly inconsistent, not that reasonable aggregation is impossible. The question is: *how much* must we sacrifice?

We provide a quantitative answer using the Latent Number ρ of the preference distribution.

1.1 Formal companion (decoder scope)

The file elysium/ fields /arrow_impossibility/arrow_impossibility_proof.py does **not** formalize full ordinal social welfare functions, profiles, or Arrow's axioms as objects in the kernel. Instead, it proves a dozen **scalar real-arithmetic lemmas** whose names mirror the narrative below. Interpreting those lemmas as social-choice theorems is an **expository bridge**: the classical impossibility logic (Arrow 1951; Gibbard 1973; Satterthwaite 1975) is standard background, while the ρ-scaling statements are **framework hypotheses** aligned with the Latent program unless explicitly noted as proved from primitives.

1.2 Contributions

- **Latent scaling narrative**: IIA stress modeled $\propto 1/\rho^2$, dictator-adjacency modeled $\propto 1/\rho$.
- **Large-N surrogate**: companion scalar inequality package compatible with gains of order $1/(N-1)$ under a linear constraint.
- **Dimension heuristic**: effective dimension read as $\log(m!)/\log(\rho)$ for m alternatives (interpretive, not the proved lemma statement).
- **Illustrated regimes**: for concentrated preferences ($\rho \geq 3$), toy calibrations suggest milder modeled violations.
- **Machine-checked companion**: 12 lemmas, no domain-specific axioms (bootstrap real arithmetic only).

2. Arrow's Axioms: Quantitative Relaxation

2.1 The Axiom Tradeoff

Theorem 1 (Sum slack for nonnegative scores). *In the companion encoding, if $d_U, d_{IIA}, d_{ND} \geq 0$ and $d_U + d_{IIA} + d_{ND} \leq B$ with $B < 3$, then $d_U + d_{IIA} + d_{ND} < 3$.*

Machine-checked: arrow_impossibility_proof.py, lemma arrow_impossibility_quantitative. □

Theorem 2 (Strict slack under a shift). *If $s_{\text{new}} = s_{\text{old}} + \Delta$ and $B < s_{\text{new}}$, then $B < s_{\text{old}} + \Delta$.*

Machine-checked: arrow_impossibility_proof.py, lemma axiom_tradeoff. □

2.2 Democratic Bound

Theorem 3 (Normalization bound). *In the companion encoding, if $N d_{ND} = N - 1$ with $N > 1$, then $d_{ND} < 1$ (equivalently $d_{ND} = (N-1)/N$ under the linear constraint).*

Machine-checked: arrow_impossibility_proof.py, lemma democratic_bound. □

3. Gibbard-Satterthwaite Extensions

3.1 Manipulation Gain

The Gibbard-Satterthwaite theorem states that any non-dictatorial, surjective voting rule for 3+ alternatives is manipulable: some voter can improve their outcome by misreporting preferences.

Theorem 4 (Strict slack from a linear per-capita constraint). *If $(N-1)g \leq 1$ with $N > 2$ and $g > 0$, then $g < 1$. For $N > 1$, the inequality $(N-1)g \leq 1$ is algebraically equivalent to $g \leq 1/(N-1)$.*

Proof sketch. The companion lemma closes $g < 1$ by nonlinear arithmetic from $(N-1)g \leq 1$ and $N > 2$. Dividing the linear constraint by $N - 1 > 0$ recovers the familiar per-capita bound $g \leq 1/(N-1)$. This is a **scalar surrogate** for "small per-capita manipulation gain" in large electorates, not a formalization of Gibbard–Satterthwaite. Machine-checked: arrow_impossibility_proof.py, lemma manipulation_gain_bound. □

3.2 Approximate Strategy-Proofness

Theorem 5 (Interior efficiency under $\eta = 1 - \varepsilon$). *If $\eta = 1 - \varepsilon$ with $0 < \varepsilon < 1$, then $0 < \eta < 1$.*

Machine-checked: arrow_impossibility_proof.py, lemma asp_efficiency_tradeoff. □

4. The Latent Spectral Framework for Social Choice

4.1 IIA Violation Frequency

The preference distribution over $\mathcal{L}^N$ has a Latent spectral decomposition with parameter ρ.

Theorem 6 (Inverse-square comparative statics in the encoding). *If $r_1^2 f_1 = r_2^2 f_2$ with $0 < r_1^2 < r_2^2$ and $f_1, f_2 > 0$, then $f_2 < f_1$.*

Narrative link. This lemma is the formal shadow of the heuristic $v(\rho) \propto C/\rho^2$: larger spectral scale ρ corresponds to smaller modeled violation intensity when holding the product $ ^2 v$ fixed across two regimes.

Machine-checked: arrow_impossibility_proof.py, lemma iia_violation_frequency. □

This is the key **interpretive** insight: IIA stress need not be uniform across preference distributions. When preferences are spectrally concentrated (high ρ), a one-dimensional latent structure may dominate, so IIA violations along subdominant directions can be comparatively rare.

4.2 Distance from Dictatorship

Theorem 7 (Linear scaling relation). *If $r\,d = C$ with $r > 1$ and $0 < C$, then $0 < d < C$.*

Narrative link. This matches the storytelling device $d_{\text{dict}}(\rho) \asymp C/\rho$ as a **proportionality template**, not a theorem about decisive voters in a concrete voting rule.

Machine-checked: arrow_impossibility_proof.py, lemma dictator_distance_from_rho. □

4.3 Social Choice Dimensionality

Theorem 8 (Relative error under scaling). *If $re_2 = e_1$ with $r > 1$ and $0 < e_2 < e_1$, then $e_2 < e_1$.*

Narrative link. The formula $d^* = \log(m!)/\log(\rho)$ remains a **designer heuristic** for effective dimension; it is not the statement proved in the companion file.

Machine-checked: arrow_impossibility_proof.py, lemma spectral_social_choice_dimension. □

4.4 Unanimity Convergence

Theorem 9 (Ordered violation intensities). *If* $r_1 v_1 = r_2 v_2$ *with* $0 < r_1 < r_2$ *and* $v_1, v_2 > 0$*, then* $v_2 < v_1$.

Narrative link. Monotone comparative statics in ρ motivate the **story** that increasing spectral concentration reduces modeled axiom stress; it does **not** assert that Arrow's three axioms become jointly satisfiable for finite $m \geq 3$ in the classical sense.

Machine-checked: arrow_impossibility_proof.py, lemma unanimity_convergence. □

5. Numerical Examples

5.1 Three Alternatives, $\rho = 3$

Metric	Value
IIA violation frequency (illustrative, $C = 1$)	$\approx 1/9 \approx 11\%$
Dictator-distance template ($d \asymp C/\rho$, $C = 1$, $\rho = 3$)	$\approx 33\%$
Heuristic dimension $\log(6)/\log(3)$ for $m = 3$	≈ 1.63
Companion linear bound ($N = 100$, $(N-1)g \leq 1$)	$g \leq 1/99 \approx 1\%$

Machine-checked numerics in arrow_impossibility_proof.py: hundred_voters_democracy (100-voter normalization + influence template) and three_alt_bound (three-alternative aggregate slack). The row $1/9 \approx 11\%$ is an **illustrative** $\rho = 3$ toy (C/ρ^2 with $C = 1$), not the equation encoded in three_alt_iia_violation (that lemma fixes a $\rho^2 = 4$ normalization, i.e. a 25% scale in the companion arithmetic). □

5.2 Interpretation

For preferences with $\rho = 3$ (moderate concentration — e.g., a society where one political dimension dominates but cross-cutting issues exist), the **toy normalization** $v \approx C/\rho^2$ with $C = 1$ suggests an $\approx 11\%$ IIA-stress rate in the Latent storytelling layer. This is **not** an empirical frequency estimate from data; it is a calibrated illustration of how spectral concentration modulates the severity of modeled violations.

In practice, this may be acceptable: most real voting systems (plurality, Borda, approval) violate IIA routinely. The Latent framework quantifies the cost and shows when it is small.

6. Discussion

6.1 Political Science Implications

The spectral framework suggests that Arrow's impossibility is most binding in societies with low ρ — i.e., where preferences are multidimensional and no dominant political dimension exists. This is consistent with the observation that political systems in highly polarized societies (effectively one-dimensional, high ρ) function "well" despite Arrow's theorem, while multi-cleavage societies struggle with aggregation.

6.2 Relationship to Existing Work

- **Kalai (2004)**: large-player robustness in games — thematically adjacent to "many voters / small per-capita influence" heuristics, distinct from Arrow's ordinal setup.
- **Mossel (2012)**: quantitative Arrow-type inequalities in a probabilistic setting — closest technical neighbor for **rates**; the present ρ-scaling is a parallel **Latent** narrative, not a reproduction of Mossel's theorem statements.
- **Procaccia & Rosenschein (2006)**: distortion metrics for voting with cardinal information — comparable **spirit** (quantifying welfare loss), without claiming extension of their formal results here.

6.3 Connection to Other Latent Economics Threads

The same ρ diagnostic is reused across other Latent economics notes (asset pricing, contagion, mechanism design, bounded rationality) as a *spectral concentration* read on whichever operator encodes uncertainty or strategic interaction in each model. This paper's social-choice use is parallel in spirit: ρ is a summary statistic for how concentrated preferences are in an abstract spectral representation, not a claim that every applied model shares identical microfoundations.

No implication is intended that Arrow's axioms become jointly satisfiable in the classical ordinal sense; the cross-domain point is methodological—one scalar for diagnosing concentration—not a unified reduction of economics to a single operator theorem.

7. Conclusion

Arrow's impossibility admits a **parametric reading** in applied work: the Latent Number ρ summarizes how spectrally concentrated preferences are, and the heuristics above suggest when modeled IIA stress and dictator-adjacency are mild. The classical impossibility for $m \geq 3$ remains; the contribution here is **diagnostic language** plus **machine-checked scalar lemmas** that keep the quantitative story honest about what is proved inside the kernel versus what is imported from economics.

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

References

- Arrow, K.J. (1951). *Social Choice and Individual Values*. Wiley.
- Gibbard, A. (1973). Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4), 587–601.
- Kalai, E. (2004). Large Robust Games. *Econometrica*, 72(6), 1631–1665.
- Mossel, E. (2012). A Quantitative Arrow Theorem. *Probability Theory and Related Fields*, 154(3–4), 735–751.
- Procaccia, A.D. & Rosenschein, J.S. (2006). The Distortion of Cardinal Preferences in Voting. *Proceedings of the 10th International Workshop on Cooperative Information Agents*, 317–331.
- Satterthwaite, M.A. (1975). Strategy-Proofness and Arrow's Conditions. *Journal of Economic Theory*, 10(2), 187–217.