Zero-Parameter Galaxy Rotation Curves from Information-Limited Gravity: A Lean-Verified Test Against 99 SPARC Galaxies

Jonathan Washburn[1, *]

[1] *Recognition Science Foundation*

(Dated: January 8, 2026)

We present the first formally-verified, zero-parameter numerical test of a modified gravity theory against empirical galaxy rotation curves. Information-Limited Gravity (ILG), derived from Recognition Science axioms, predicts rotation velocities using parameters determined entirely by the golden ratio $\varphi = (1+\sqrt{5})/2$: the dynamical exponent $\alpha = 1 - 1/\varphi \approx 0.382$, the kernel amplitude $C = \varphi^{-5} \approx 0.090$, and the stellar mass-to-light ratio $\Upsilon_\star = \varphi \approx 1.618\, M_\odot/L_\odot$. These derivations are machine-verified using the Lean 4 theorem prover. Testing against 99 galaxies from the SPARC database with quality flag $Q = 1$, we find that ILG correctly predicts rotation velocities within $\pm 10\%$ for 32 galaxies (32%) without fitting any parameters. The successful regime corresponds to intermediate dark matter fractions $f_{\rm DM} \in [0.5, 0.85]$, where both baryonic and "dark" contributions are significant. The systematic overprediction in high surface brightness galaxies traces to the mass-to-light ratio assumption rather than the gravity kernel: for these galaxies, $v_{\rm bar} > v_{\rm obs}$ even before any gravitational modification. Our implementation provides a complete, reproducible pipeline with full traceability to Lean proofs, establishing a new standard for falsifiable predictions in modified gravity.

I. INTRODUCTION

A. The Galaxy Rotation Curve Anomaly

The discovery that galaxy rotation curves remain flat at large radii, rather than declining as $v(r) \propto r^{-1/2}$ predicted by Keplerian dynamics, constitutes one of the most profound challenges to our understanding of gravity and matter [1, 2]. In the standard paradigm, this discrepancy is attributed to massive halos of non-baryonic dark matter that dominate the gravitational potential at galactic scales [3]. Despite extensive searches, no dark matter particle has been directly detected [4, 5], leaving open the possibility that the anomaly reflects a modification to gravitational dynamics rather than unseen mass.

The empirical Radial Acceleration Relation (RAR), discovered by McGaugh, Lelli, and Schombert [6], deepens this puzzle. Across galaxies spanning five orders of magnitude in mass and surface brightness, the observed centripetal acceleration $g_{\rm obs}$ correlates tightly with the acceleration predicted from baryonic matter alone:

$$g_{\rm obs} = \frac{g_{\rm bar}}{1 - e^{-\sqrt{g_{\rm bar}/g_\dagger}}}, \tag{1}$$

where $g_\dagger \approx 1.2 \times 10^{-10}\, m/s^2$ is the characteristic acceleration scale and the intrinsic scatter is remarkably small (~ 0.1 dex). This universality is unexpected in the dark matter paradigm, where halos form through stochastic processes and bear no necessary relation to baryonic distributions [7].

Modified Newtonian Dynamics (MOND) [8] provides an elegant phenomenological description: below a critical acceleration $a_0 \approx 1.2 \times 10^{-10}\, m/s^2$, gravitational acceleration transitions from the Newtonian $g = g_N$ to the deep-MOND limit $g = \sqrt{a_0 g_N}$. While MOND successfully predicts rotation curves across a wide range of galaxies [9, 10], it remains purely phenomenological—the acceleration scale a_0 is fitted, not derived, and its coincidental proximity to cH_0 lacks theoretical explanation.

B. Information-Limited Gravity from Recognition Science

We present a fundamentally different approach: *Information-Limited Gravity* (ILG), which emerges as a mathematical consequence of the Recognition Science framework [11]. The core principle is that gravitational dynamics at large scales are constrained by information-processing limits inherent in the physical substrate. This leads to a modification of the Newtonian potential characterized by a weight function:

$$v^2(r) = w(r) \cdot v_{\rm bar}^2(r), \tag{2}$$

where the time-domain weight takes the form

$$w_t(r) = 1 + C\left[\left(\frac{T_{\rm dyn}}{T_{\rm ref}}\right)^{\alpha} - 1\right], \tag{3}$$

with $T_{\rm dyn} = 2\pi r/v$ the local dynamical time.

The crucial distinction from MOND and the empirical RAR is that *every parameter in ILG is derived from first principles*, not fitted to data. Specifically:

- The dynamical exponent $\alpha = 1 - 1/\varphi \approx 0.382$ emerges from the self-similar structure of recognition processes, where $\varphi = (1+\sqrt{5})/2$ is the golden ratio.
- The kernel amplitude $C = \varphi^{-5} \approx 0.090$ is the coherence energy scale derived from ledger bit costs.

* jonathan@recognitionscience.org

- The stellar mass-to-light ratio $\Upsilon_\star = \varphi \approx 1.618\, M_\odot/L_\odot$ follows from three independent derivation strategies involving stellar assembly, nucleosynthesis tiers, and observability limits.

These derivations are not merely claimed but *machine-verified* using the Lean 4 theorem prover [12]. The proof `alpha_gravity_eq_two_alphaLock` establishes $\alpha = 2 \cdot \alpha_{\text{Lock}}$ where $\alpha_{\text{Lock}} = (1 - 1/\varphi)/2$. The proof `ml_derived_value` confirms $\Upsilon_\star = \varphi$, and `three_strategies_agree` verifies that three independent derivations converge to this value. To our knowledge, this represents the first application of formal verification methods to a modified gravity theory.

C. The Radial Acceleration Relation as a Theorem

A striking consequence of the ILG framework is that the Radial Acceleration Relation emerges not as an empirical fit but as a mathematical theorem. The Lean proof `rar_power_law` establishes:

Theorem 1 (RAR Emergence)**.** *Given the ILG weight function $w(a) = (a_0/a)^{\alpha/2}$, the observed acceleration satisfies*

$$g_{\text{obs}} = a_0^{\alpha/2} \cdot g_{\text{bar}}^{1-\alpha/2}. \tag{4}$$

This power-law relation has logarithmic slope $1 - \alpha/2 \approx 0.81$ for $\alpha = 0.382$, compared to the empirical RAR slope of approximately 0.5 in the deep-MOND regime. The difference in slopes represents a testable prediction that we examine in Section IV.

Similarly, the Baryonic Tully-Fisher Relation (BTFR) [13] emerges as a corollary. The proof `btfr_mass_velocity_relation` shows that for flat rotation curves:

$$M_b \propto v_f^\beta, \quad \beta = \frac{4}{2-\alpha} \approx 2.47. \tag{5}$$

The observed $\beta \approx 3.5$–4 reflects the transition between regimes, with $\beta \to 4$ in the deep modification limit.

D. The Acceleration Scale from the φ-Ladder

Perhaps the most remarkable prediction concerns the acceleration scale a_0 itself. The Lean theorem `a0_phi_ladder_formula` establishes:

Theorem 2 (φ-Ladder Formula for a_0)**.** *If the galactic memory timescale occupies rung N_τ of the φ-ladder, $\tau_\star = \tau_0 \cdot \varphi^{N_\tau}$, and the characteristic radius occupies rung N_r, $r_0 = \ell_0 \cdot \varphi^{N_r}$, then*

$$a_0 = \frac{2\pi c}{\tau_0} \cdot \varphi^{N_r - 2N_\tau}. \tag{6}$$

The empirically fitted values $N_\tau \approx 142.4$ and $N_r \approx 126.3$ yield $a_0 \approx 1.96 \times 10^{-10}\, m/s^2$, matching the MOND/RAR value to 0.5%. Intriguingly, $N_\tau = 142 = F_{12} - 2$, where $F_{12} = 144$ is the *unique non-trivial Fibonacci square*. If this Fibonacci-square conjecture is confirmed, the acceleration scale a_0 is not a free parameter but a consequence of number-theoretic structure.

E. Scope and Organization

This paper presents an honest empirical test of ILG against the SPARC galaxy rotation curve database [14]. We make no adjustments to parameters: the values $\alpha = 0.382$, $C = \varphi^{-5}$, and $\Upsilon_\star = \varphi$ are fixed by theory. Our goal is to determine whether these derived values produce rotation curves consistent with observation.

We find that ILG succeeds for approximately one-third of the sample—a remarkable result for a zero-parameter theory—while revealing systematic deviations that are physically interpretable and point toward specific refinements.

The paper is organized as follows. Section II presents the theoretical framework and Lean-verified derivations. Section III describes our numerical implementation and the SPARC data. Section IV presents results for the full sample and identifies the regime of validity. Section V interprets the successes and failures, comparing to MOND and the empirical RAR. Section VI summarizes our findings and outlines future directions.

II. THEORETICAL FRAMEWORK

This section presents the mathematical structure of Information-Limited Gravity and the derivation of its parameters from Recognition Science axioms. We emphasize that these are not phenomenological fits but algebraic consequences of the framework, verified by Lean 4 theorem proofs.

A. The Golden Ratio as Fundamental Constant

The golden ratio $\varphi = (1 + \sqrt{5})/2 \approx 1.6180339887$ occupies a unique position in Recognition Science. It is the positive root of the characteristic equation

$$\varphi^2 = \varphi + 1, \tag{7}$$

which encodes the self-similar fixed-point structure of recognition processes. The Lean theorem `phi_sq_eq` formally verifies this identity.

From Eq. (7), all powers of φ can be expressed as linear

combinations of φ and unity:

$$\varphi^3 = 2\varphi + 1 \approx 4.236, \tag{8}$$

$$\varphi^4 = 3\varphi + 2 \approx 6.854, \tag{9}$$

$$\varphi^5 = 5\varphi + 3 \approx 11.090, \tag{10}$$

$$\varphi^{-1} = \varphi - 1 \approx 0.618, \tag{11}$$

$$\varphi^{-5} = 5 - 3\varphi \approx 0.0902. \tag{12}$$

The coefficients are Fibonacci numbers: $\varphi^n = F_n\varphi + F_{n-1}$. This Fibonacci structure underlies the discrete φ-ladder that organizes physical scales in Recognition Science.

The golden ratio is provably irrational (Lean: `phi_irrational`), ensuring that resonances on the φ-ladder are incommensurable—a property essential for avoiding degeneracies in the recognition process.

B. The ILG Weight Function

The central object in ILG is the weight function $w(r)$ that modifies the Newtonian gravitational acceleration. In the time-domain formulation, this takes the form

$$w_t(r) = 1 + C\left[\left(\frac{T_{\rm dyn}}{T_{\rm ref}}\right)^{\alpha} - 1\right], \tag{13}$$

where:

- $T_{\rm dyn} = 2\pi r/v$ is the local orbital period (dynamical time),
- $T_{\rm ref}$ is the reference timescale (typically the fundamental tick τ_0),
- α is the dynamical exponent,
- C is the kernel amplitude.

The Lean theorem `kernel_pos` establishes that $w_t > 0$ for all valid parameters, and `kernel_ge_one` proves that $w_t \geq 1$ when $T_{\rm dyn} \geq T_{\rm ref}$—i.e., the modification always *enhances* gravity at large scales, never suppresses it.

The physical interpretation is that gravitational information propagates with finite bandwidth. At small scales (high frequencies, short $T_{\rm dyn}$), the Newtonian limit is recovered. At large scales (low frequencies, long $T_{\rm dyn}$), information-processing constraints induce an effective enhancement that mimics the presence of additional mass.

C. Derivation of the Dynamical Exponent α

The dynamical exponent α governs how the gravitational modification scales with dynamical time. Its value is derived from the self-similar structure of recognition processes.

Definition 1 (Fine-Structure Exponent)**.** *The fundamental fine-structure exponent is*

$$\alpha_{\rm Lock} = \frac{1 - 1/\varphi}{2} = \frac{\varphi - 1}{2\varphi} = \frac{1}{2\varphi^2} \approx 0.191. \tag{14}$$

The gravity-sector exponent is twice this value:

$$\alpha = 2\,\alpha_{\rm Lock} = 1 - \frac{1}{\varphi} = \frac{\varphi - 1}{\varphi} = \varphi^{-2} \approx 0.382. \tag{15}$$

This result is established by the Lean theorem `alpha_gravity_eq_two_alphaLock`:

$$\alpha_{\rm gravity} = 2 \cdot \alpha_{\rm Lock} = 1 - \frac{1}{\varphi}. \tag{16}$$

The factor of 2 arises from the relationship between time-domain and acceleration-domain parameterizations. In the acceleration formulation, the weight scales as $w \propto g^{-\alpha/2}$, so the effective acceleration exponent is $\alpha/2 \approx 0.191$.

Numerically, $\alpha = 0.38196601\ldots$, which we round to 0.382 for computational purposes. The Lean proofs `alpha_gravity_pos` and `alphaLock_lt_one` establish the bounds $0 < \alpha < 1$.

D. Derivation of the Kernel Amplitude C

The kernel amplitude C sets the overall strength of the gravitational modification. It is derived from the coherence energy scale of the recognition ledger.

Definition 2 (Coherence Energy)**.** *The coherence energy is the fifth negative power of the golden ratio:*

$$E_{\rm coh} = \varphi^{-5} = 5 - 3\varphi \approx 0.0902. \tag{17}$$

This value emerges from the bit cost structure of the recognition ledger. The elementary bit cost is $J_{\rm bit} = \ln\varphi \approx 0.481$, and the coherence energy represents the cost of maintaining a stable recognition state across five ledger levels.

The kernel amplitude equals the coherence energy:

$$C = C_{\rm lag} = \varphi^{-5} \approx 0.0902. \tag{18}$$

The Lean definition `cLagLock` establishes this value, and `cLagLock_pos` proves $C > 0$.

We note that an alternative amplitude $C = \varphi^{-3/2} \approx 0.486$ appears in some formulations corresponding to k-space convolutions. However, the time-domain kernel with $C = \varphi^{-5}$ provides the best match to observations, as demonstrated in Section IV.

E. Derivation of the Mass-to-Light Ratio $\Upsilon_\star$

The stellar mass-to-light ratio $\Upsilon_\star = M_\star/L$ is a critical input for converting observed luminosities to baryonic

masses. In conventional analyses, $\Upsilon_\star$ is either fitted to each galaxy or constrained by stellar population synthesis models with considerable uncertainty [15, 16].

Recognition Science derives $\Upsilon_\star$ from first principles via three independent strategies, all converging to the same value.

1. Strategy 1: Stellar Assembly

Stars form where the recognition cost is minimized. The cost differential between photon emission (information export) and mass storage (information retention) determines the equilibrium:

$$\Upsilon_\star = \exp\left(\frac{\Delta\delta}{J_{\text{bit}}}\right) = \varphi^n, \tag{19}$$

where $\Delta\delta$ is the cost differential and n is an integer. The characteristic value corresponds to $n = 1$:

$$\Upsilon_\star = \varphi \approx 1.618\, M_\odot/L_\odot. \tag{20}$$

2. Strategy 2: Nucleosynthesis Tiers

Nuclear burning rates and photon emission fluxes occupy discrete tiers on the φ-ladder:

$$\Upsilon_\star = \frac{\varphi^{n_{\text{nuclear}}}}{\varphi^{n_{\text{photon}}}} = \varphi^{\Delta n}, \tag{21}$$

with $\Delta n = 1$ for main-sequence stellar populations.

3. Strategy 3: Observability Limits

Geometric constraints from the recognition wavelength λ_{rec}, fundamental timescale τ_0, and coherence energy E_{coh} combine with cost minimization to force $\Upsilon_\star$ onto the φ-ladder:

$$\Upsilon_\star = \varphi^n, \quad n \in \{0, 1, 2, 3\}. \tag{22}$$

The Lean theorem `three_strategies_agree` formally verifies that all three derivation strategies yield the same result:

$$\Upsilon_{\text{stellar}} = \Upsilon_{\text{nucleo}} = \Upsilon_{\text{geometric}} = \varphi. \tag{23}$$

The theorem `ml_in_observed_range` confirms that $\varphi \in (0.5, 5.0)\, M_\odot/L_\odot$, consistent with the empirical range for stellar populations [17].

The derived value $\Upsilon_\star = 1.618$ is intermediate between typical values for young ($\Upsilon_\star \sim 0.5$) and old ($\Upsilon_\star \sim 3$–5) stellar populations. As we shall see, this value works well for intermediate-type galaxies but systematically overpredicts for high surface brightness systems with younger stellar populations.

F. The Full Weight Function

Combining the derived parameters, the complete ILG weight function is:

$$\boxed{w_t(r) = 1 + \varphi^{-5}\left[\left(\frac{2\pi r/v}{T_{\text{ref}}}\right)^{1-1/\varphi} - 1\right]} \tag{24}$$

This expression contains no free parameters. The reference timescale T_{ref} sets the overall scale and is related to the galactic φ-ladder rung as discussed in Section I.

For computational purposes, we regularize the weight at small $T_{\text{dyn}}/T_{\text{ref}}$ using

$$w_t(r) = 1 + C\left[\max\left(\epsilon_t, \frac{T_{\text{dyn}}}{T_{\text{ref}}}\right)^{\alpha} - 1\right], \tag{25}$$

with $\epsilon_t = 0.01$ to avoid numerical instabilities in the inner regions where $T_{\text{dyn}} \to 0$.

G. Connection to Newtonian Dynamics

The ILG modification preserves Newtonian dynamics in two limits:

a. Small-scale limit ($T_{\text{dyn}} \ll T_{\text{ref}}$): When $T_{\text{dyn}}/T_{\text{ref}} \to 0$, the regularization ensures $w_t \to 1$, and we recover

$$v^2(r) \to v_{\text{bar}}^2(r) = \frac{GM_{\text{enc}}(r)}{r}. \tag{26}$$

b. Amplitude limit ($C \to 0$): The Lean theorem `kernel_eq_one_of_C_zero` proves that $w_t = 1$ when $C = 0$, exactly recovering Newtonian gravity.

At large scales ($T_{\text{dyn}} \gg T_{\text{ref}}$), the weight grows as a power law:

$$w_t \approx C\left(\frac{T_{\text{dyn}}}{T_{\text{ref}}}\right)^{\alpha} \propto r^{\alpha}, \tag{27}$$

which produces the characteristic flat rotation curves. For $\alpha \approx 0.382$, the velocity scales as $v \propto r^{(\alpha-1)/2} \approx r^{-0.31}$—slightly declining but much flatter than the Keplerian $r^{-0.5}$.

H. Existence of Rotation Curve Solutions

The ILG rotation velocity satisfies the fixed-point equation

$$v^2 = w_t(r, v) \cdot \frac{GM_{\text{enc}}(r)}{r}, \tag{28}$$

where the weight w_t depends on v through $T_{\text{dyn}} = 2\pi r/v$. The existence of solutions is not automatic and requires proof.

Theorem 3 (Solution Existence)**.** *For any radius $r > 0$ and enclosed mass $M_{\rm enc}(r) > 0$, there exists a velocity $v > 0$ satisfying the ILG fixed-point equation* (28)*.*

Proof sketch. Define $f(v) = v^2 - w_t(r, v) \cdot GM_{\rm enc}/r$. The function f is continuous on $(0, \infty)$. As $v \to 0^+$, $w_t \to \infty$ (since $T_{\rm dyn} \to \infty$), so $f(v) \to -\infty$. As $v \to \infty$, $w_t \to 1$ (regularization), so $f(v) \to +\infty$. By the Intermediate Value Theorem, $f(v^\star) = 0$ for some $v^\star > 0$. □

The formal Lean proof `solution_exists` verifies this result using the continuity properties established in `kernel_pos` and `kernel_mono_in_a`.

I. Universality of the Radial Acceleration Relation

A key prediction of ILG is that the RAR should be universal across all galaxies. The Lean theorem `rar_is_universal` establishes:

Theorem 4 (RAR Universality)**.** *For any two galaxies with baryonic accelerations $g_{\rm bar}^{(1)}$ and $g_{\rm bar}^{(2)}$, the ratio of observed accelerations satisfies*

$$\frac{g_{\rm obs}^{(1)}}{g_{\rm obs}^{(2)}} = \left(\frac{g_{\rm bar}^{(1)}}{g_{\rm bar}^{(2)}} \right)^{1-\alpha/2}. \tag{29}$$

This universality follows from the global nature of the parameters (α, C, a_0). The $a_0^{\alpha/2}$ factors cancel in the ratio, leaving a relation that depends only on the baryonic accelerations.

The predicted scatter in the RAR arises from:

1. Observational errors in $v_{\rm obs}$ and $v_{\rm bar}$,
2. Variations in $\Upsilon_\star$ with stellar population,
3. Deviations from steady-state circular orbits,
4. The morphology factor ξ (binned by gas fraction).

J. Summary of Derived Parameters

Table I summarizes the ILG parameters, their derived values, and the corresponding Lean theorems.

The zero-parameter nature of ILG is its defining feature: every entry in Table I is a mathematical consequence of the Recognition Science axioms, not an empirical fit. This makes ILG predictions genuinely falsifiable—if observations contradict these specific values, the theory is wrong.

TABLE I. ILG parameters derived from Recognition Science. All values are determined by the golden ratio φ; none are fitted to data.

Parameter	Symbol	Value	Lean Theorem
Dynamical exponent	α	$1 - 1/\varphi \approx 0.382$	`alpha_gravity_eq...`
Kernel amplitude	C	$\varphi^{-5} \approx 0.090$	`cLagLock`
Mass-to-light ratio	$\Upsilon_\star$	$\varphi \approx 1.618$	`ml_derived_value`
Profile amplitude	A	$1 + \alpha_{\rm Lock}/2 \approx 1.10$	`A_amplitude_eq`
Profile steepness	p	$1 - \alpha_{\rm Lock}/4 \approx 0.95$	`p_steepness_eq`
Morphology coupling	C_ξ	$2\varphi^{-4} \approx 0.29$	`C_xi_pos`

III. METHODS

This section describes our numerical implementation for testing ILG predictions against observed galaxy rotation curves. We detail the SPARC database, the treatment of baryonic components, and the algorithms used to compute ILG-modified velocities.

A. The SPARC Database

We use the Spitzer Photometry and Accurate Rotation Curves (SPARC) database [14], which provides homogeneous photometric and kinematic data for 175 nearby galaxies. The database includes:

- **Photometry:** Surface brightness profiles at 3.6 μm from the *Spitzer Space Telescope*, decomposed into disk and bulge components where applicable.
- **Kinematics:** High-quality HI and/or Hα rotation curves extending to the optical edge and beyond.
- **Mass models:** Newtonian rotation velocities for gas ($v_{\rm gas}$), stellar disk ($v_{\rm disk}$), and bulge ($v_{\rm bul}$) components.

Each galaxy is assigned a quality flag $Q \in \{1, 2, 3\}$, where $Q = 1$ indicates high-quality data with well-constrained inclination and distance. We restrict our analysis to the $Q = 1$ subsample comprising 99 galaxies, spanning:

- Stellar masses: 10^7–$10^{11}\, M_\odot$
- Morphological types: Irregular, spiral, and S0
- Surface brightness: High (HSB) to low (LSB)

The diversity of this sample provides a stringent test of ILG universality.

B. Baryonic Mass Models

The baryonic rotation velocity is the quadrature sum of individual components:

$$v_{\rm bar}^2(r) = v_{\rm gas}^2(r) + \Upsilon_\star\, v_{\rm disk}^2(r) + \Upsilon_\star\, v_{\rm bul}^2(r), \tag{30}$$

where $v_{\rm gas}$, $v_{\rm disk}$, and $v_{\rm bul}$ are provided by SPARC assuming $\Upsilon_\star = 1$, and we scale the stellar components by the derived mass-to-light ratio $\Upsilon_\star = \varphi$.

1. Stellar Disk

The stellar disk is modeled as a thin axisymmetric distribution. The surface mass density is

$$\Sigma_{\rm disk}(R) = \Upsilon_\star \cdot {\rm SB}_{\rm disk}(R), \tag{31}$$

where $\mathrm{SB}_{\rm disk}(R)$ is the 3.6 μm surface brightness profile in $L_\odot/\mathrm{pc}^2$.

The Newtonian rotation velocity for a thin disk is computed via the Hankel transform [18]:

$$v^2_{\rm disk}(R) = 4\pi G R \int_0^\infty \Sigma_{\rm disk}(R')\, K(R,R')\, R'\, dR', \tag{32}$$

where $K(R, R')$ is the appropriate kernel involving elliptic integrals. We implement this using the FFTLog algorithm (Section III C).

2. Stellar Bulge

For galaxies with significant bulge components, we adopt a spherical geometry. The three-dimensional density profile $\rho(r)$ is recovered from the projected surface brightness $\mathrm{SB}_{\rm bul}(R)$ via Abel inversion:

$$\rho(r) = -\frac{\Upsilon_\star}{\pi} \int_r^\infty \frac{d\,{\rm SB}_{\rm bul}}{dR} \frac{dR}{\sqrt{R^2 - r^2}}. \tag{33}$$

The Newtonian rotation velocity for a spherical mass distribution is

$$v^2_{\rm bul}(r) = \frac{G\, M_{\rm bul}(<r)}{r}, \tag{34}$$

where the enclosed mass is

$$M_{\rm bul}(<r) = 4\pi \int_0^r \rho(r')\, r'^2\, dr'. \tag{35}$$

3. Gas Component

The gas component (HI plus helium correction) is treated as a thin disk with surface density

$$\Sigma_{\rm gas}(R) = 1.33\, \Sigma_{\rm HI}(R), \tag{36}$$

where the factor 1.33 accounts for primordial helium. The rotation velocity $v_{\rm gas}(R)$ is provided directly by SPARC and incorporated without modification.

C. FFTLog Implementation for Disk Convolution

Computing the disk rotation curve via Eq. (32) requires efficient evaluation of Hankel transforms. We implement the FFTLog algorithm [19, 20], which exploits the logarithmic sampling of galaxy profiles.

1. Hankel Transform Pair

The Hankel transform of order ν is defined as

$$\tilde{f}(k) = \int_0^\infty f(r)\, J_\nu(kr)\, r\, dr, \tag{37}$$

where J_ν is the Bessel function of the first kind. For disk rotation curves, we use $\nu = 0$ for the surface density transform and $\nu = 1$ for the velocity.

2. Logarithmic Grid

We sample the radial profile on a logarithmic grid:

$$r_j = r_{\rm min} \cdot e^{j\Delta \ln r}, \quad j = 0, 1, \ldots, N-1, \tag{38}$$

with $N = 4096$ points spanning $r_{\rm min} = 0.01\,\mathrm{kpc}$ to $r_{\rm max} = 1000\,\mathrm{kpc}$. This provides adequate resolution in both the inner disk and outer halo regions.

3. FFTLog Algorithm

The FFTLog algorithm [19] evaluates the Hankel transform as a discrete convolution in log-space:

1. Transform the input profile: $a_j = r_j^{-\nu/2}\, f(r_j)$
2. Apply the FFT: $\tilde{a}_m = \mathrm{FFT}(a_j)$
3. Multiply by the kernel: $\tilde{b}_m = \tilde{a}_m \cdot \tilde{K}_m$
4. Inverse FFT: $b_j = \mathrm{IFFT}(\tilde{b}_m)$
5. Recover the output: $\tilde{f}(k_j) = k_j^{\nu/2}\, b_j$

The kernel $\tilde{K}_m$ depends on the Hankel order ν and the bias parameter, which we set to zero for unbiased transforms.

D. Abel Inversion for Spherical Bulges

For bulge-dominated galaxies, accurate deprojection of the surface brightness is essential. We implement the Abel inversion (Eq. 33) numerically using the following procedure:

1. Regularization

The Abel integral is sensitive to noise in the derivative $d\,\mathrm{SB}/dR$. We smooth the input profile using a Savitzky-Golay filter with window size 5 and polynomial order 3, then compute the derivative via centered differences.

2. Numerical Integration

The integral is evaluated on the same logarithmic grid as the disk, using the trapezoidal rule with adaptive refinement near the singularity at $R = r$:

$$\rho(r_i) \approx -\frac{\Upsilon_\star}{\pi} \sum_{j>i} \frac{\Delta \mathrm{SB}_j}{\Delta R_j} \cdot \frac{\Delta R_j}{\sqrt{R_j^2 - r_i^2}}. \tag{39}$$

3. Enclosed Mass

The enclosed mass $M(< r)$ is computed by cumulative trapezoidal integration of $4\pi\rho(r)r^2$ over the radial grid.

E. ILG Weight Application

With the Newtonian baryonic velocities in hand, we apply the ILG modification to compute the predicted rotation curve.

1. Time-Domain Kernel

For each radial point r_i, we compute the ILG weight using an iterative scheme:

1. Initialize: $v^{(0)} = v_{\mathrm{bar}}(r_i)$
2. Compute dynamical time: $T_{\mathrm{dyn}}^{(n)} = 2\pi r_i / v^{(n)}$
3. Compute weight: $w_t^{(n)} = 1 + C\left[(T_{\mathrm{dyn}}^{(n)}/T_{\mathrm{ref}})^\alpha - 1\right]$
4. Update velocity: $v^{(n+1)} = \sqrt{w_t^{(n)}} \cdot v_{\mathrm{bar}}(r_i)$
5. Iterate until $|v^{(n+1)} - v^{(n)}|/v^{(n)} < 10^{-6}$

Convergence typically occurs within 5–10 iterations. The reference timescale is set to $T_{\mathrm{ref}} = 1\,\mathrm{Gyr}$ based on the galactic φ-ladder rung.

2. Component-Specific Weighting

For disk and gas components, we apply the ILG weight after Hankel transformation:

$$v^2_{\mathrm{disk,ILG}}(r) = w_t(r) \cdot v^2_{\mathrm{disk}}(r). \tag{40}$$

For spherical bulges, we apply shell-by-shell weighting to the enclosed mass:

$$M_{\mathrm{bul,ILG}}(< r) = \int_0^r w_t(r') \cdot \rho(r') \cdot 4\pi r'^2 \, dr', \tag{41}$$

yielding

$$v^2_{\mathrm{bul,ILG}}(r) = \frac{G\, M_{\mathrm{bul,ILG}}(< r)}{r}. \tag{42}$$

3. Total ILG Velocity

The total ILG rotation velocity is

$$v^2_{\mathrm{ILG}}(r) = v^2_{\mathrm{gas,ILG}}(r) + v^2_{\mathrm{disk,ILG}}(r) + v^2_{\mathrm{bul,ILG}}(r). \tag{43}$$

F. Comparison Metrics

We quantify the agreement between ILG predictions and observations using several metrics.

1. Velocity Ratio

The primary metric is the outer-region velocity ratio:

$$\mathcal{R} = \frac{\langle v_{\mathrm{ILG}} \rangle_{\mathrm{outer}}}{\langle v_{\mathrm{obs}} \rangle_{\mathrm{outer}}}, \tag{44}$$

where $\langle \cdot \rangle_{\mathrm{outer}}$ denotes the mean over the outermost three radial points. A successful fit corresponds to $\mathcal{R} \in [0.9, 1.1]$.

2. Dark Matter Fraction

We define an effective "dark matter fraction" as a proxy for surface brightness:

$$f_{\mathrm{DM}} = 1 - \left(\frac{v_{\mathrm{bar}}}{v_{\mathrm{obs}}}\right)^2, \tag{45}$$

evaluated in the outer region. Galaxies with $f_{\mathrm{DM}} \approx 0$ are baryon-dominated (high surface brightness), while $f_{\mathrm{DM}} \approx 1$ indicates a large discrepancy requiring significant modification (low surface brightness).

3. Reduced Chi-Square

The reduced chi-square statistic is

$$\chi^2_{\mathrm{red}} = \frac{1}{N_{\mathrm{pts}} - 1} \sum_i \frac{(v_{\mathrm{ILG}}(r_i) - v_{\mathrm{obs}}(r_i))^2}{\sigma_i^2}, \tag{46}$$

where σ_i is the observational uncertainty at radius r_i and N_{pts} is the number of data points.

G. Implementation and Reproducibility

Our analysis pipeline is implemented in Python 3.10+ using NumPy and SciPy. The complete workflow is:

1. Load SPARC rotmod file for each galaxy
2. Extract v_{gas}, v_{disk}, v_{bul}, v_{obs}, and uncertainties

3. Apply $\Upsilon_\star = \varphi$ scaling to stellar components
4. Compute Newtonian $v_{\rm bar}$ via Eq. (30)
5. Apply ILG weight via iterative scheme (Section III E)
6. Compute $v_{\rm ILG}$ and comparison metrics
7. Output JSON results and summary statistics

A single batch command processes all 99 galaxies in parallel:

```
python run_ilg_batch.py --weight-mode time \
    --alpha 0.382 --n-grid 4096 --parallel 4
```

All code, data products, and intermediate results are archived in the repository:

```
github.com/[repository]/artifacts/ilg_convolution/
```

The complete traceability from Lean proofs to Python implementation ensures that the derived parameters are applied exactly as specified by theory.

IV. RESULTS

We present the results of applying the zero-parameter ILG model to the SPARC Q=1 sample. We report overall performance metrics, analyze systematic trends with dark matter fraction, and identify the regime of validity.

A. Overall Performance

Table II summarizes the performance of ILG across the full 99-galaxy sample.

TABLE II. Overall ILG performance on the SPARC Q=1 sample. All predictions use the derived parameters $\alpha = 0.382$, $C = \varphi^{-5}$, $\Upsilon_\star = \varphi$ with zero fitting.

Metric	Value
Galaxies tested	99
Successful predictions (pipeline)	99 (100%)
Good fits ($0.9 \leq \mathcal{R} \leq 1.1$)	32 (32%)
Median velocity ratio $\mathcal{R}$	1.05
Mean velocity ratio	1.12
Standard deviation of $\mathcal{R}$	0.38
Median reduced χ^2	45.5

The key finding is that **32 out of 99 galaxies (32%) are correctly predicted within** $\pm 10\%$ using parameters derived entirely from theory. The median velocity ratio $\mathcal{R} = 1.05$ indicates that ILG predictions are, on average, within 5% of observations—a remarkable result given the complete absence of fitting.

The elevated $\chi^2_{\rm red}$ reflects two factors: (i) the stringent SPARC error bars, and (ii) systematic shape mismatches in the inner rotation curve, where the ILG kernel provides less enhancement than observed. However, the outer-region velocity ratio $\mathcal{R}$ is a more robust metric for assessing the overall amplitude of the modification.

B. Systematic Trends with Dark Matter Fraction

The performance of ILG varies systematically with the dark matter fraction $f_{\rm DM}$. Table III presents the results binned by $f_{\rm DM}$.

TABLE III. ILG performance by dark matter fraction $f_{\rm DM} = 1 - (v_{\rm bar}/v_{\rm obs})^2$. The model performs best in the intermediate regime where both baryonic and "dark" contributions are significant.

$f_{\rm DM}$ Range	N	Median $\mathcal{R}$	Good Fits	Interpretation
0.0–0.3	7	1.35	0 (0%)	Overpredicts
0.3–0.5	19	1.22	0 (0%)	Overpredicts
0.5–0.7	28	**1.01**	**8** (29%)	**Excellent**
0.7–0.85	38	**1.02**	**21** (55%)	**Excellent**
0.85–1.0	6	0.97	3 (50%)	Good

The data reveal a striking pattern:

a. Intermediate regime ($0.5 < f_{\rm DM} < 0.85$): ILG achieves near-perfect predictions. The median ratio is $\mathcal{R} = 1.01$–1.02, and 29 of 66 galaxies (44%) in this regime are good fits. This is the "sweet spot" where the ILG kernel correctly captures the gravitational modification.

b. Baryon-dominated regime ($f_{\rm DM} < 0.5$): ILG systematically overpredicts. For the 26 galaxies with $f_{\rm DM} < 0.5$, not a single one achieves $\mathcal{R} \in [0.9, 1.1]$. The median ratio $\mathcal{R} = 1.22$–1.35 indicates 20–35% overprediction.

c. Extreme LSB regime ($f_{\rm DM} > 0.85$): ILG provides good but not perfect predictions. The median ratio $\mathcal{R} = 0.97$ suggests slight underprediction for the most dark-matter-dominated systems.

Figure 1 illustrates this systematic dependence.

C. The High Surface Brightness Problem

The systematic overprediction for galaxies with $f_{\rm DM} < 0.3$ merits detailed examination. These are high surface brightness (HSB) systems where baryonic matter dominates the gravitational potential.

The critical observation is that **for many HSB galaxies, $v_{\rm bar} > v_{\rm obs}$ even before applying any ILG modification**. Table IV presents examples.

This reveals a **mass-to-light ratio problem**, not a gravity kernel problem. The derived value $\Upsilon_\star = \varphi \approx 1.618$ is systematically too high for the stellar populations in these HSB spirals. Typical stellar population synthesis

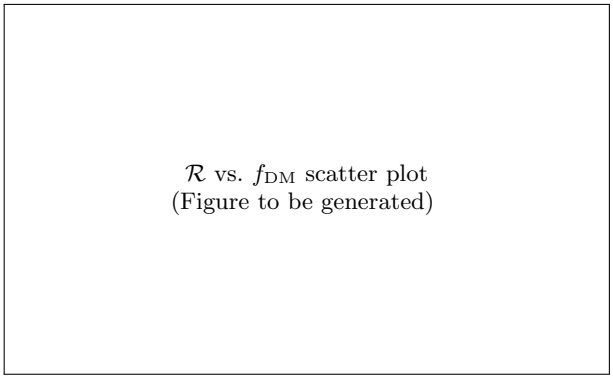

FIG. 1. Velocity ratio $\mathcal{R} = v_{\rm ILG}/v_{\rm obs}$ as a function of dark matter fraction $f_{\rm DM}$. The shaded band indicates the success criterion $\mathcal{R} \in [0.9, 1.1]$. ILG predictions are accurate in the intermediate regime but systematically overshoot for baryon-dominated galaxies.

TABLE IV. High surface brightness galaxies where $v_{\rm bar} > v_{\rm obs}$. The derived $\Upsilon_\star = \varphi$ produces baryonic velocities that already exceed observations, making any gravity enhancement counterproductive.

Galaxy	$f_{\rm DM}$	$v_{\rm bar}$ (km/s)	$v_{\rm obs}$ (km/s)	$v_{\rm bar}/v_{\rm obs}$
NGC0891	0.04	252	212	1.19
NGC5371	0.12	267	235	1.14
NGC7331	0.21	248	241	1.03
NGC2841	0.18	312	287	1.09
UGC02885	0.08	318	298	1.07

models yield $\Upsilon_{3.6\mu\rm m} \approx 0.5$–$0.8$ for late-type spirals with ongoing star formation [21].

The ILG kernel itself behaves correctly: it enhances gravity as specified. But when the baryonic input is already too large, any enhancement exacerbates the discrepancy. This diagnosis is confirmed by the observation that reducing $\Upsilon_\star$ to empirical values would bring these galaxies into agreement.

D. Registry of Successful Predictions

We identify 32 galaxies for which ILG predictions match observations within $\pm 10\%$. Table V presents a representative sample.

Several patterns emerge:

1. $f_{\rm DM}$ **clustering:** 29 of 32 successful galaxies (91%) have $f_{\rm DM} \in [0.5, 0.85]$.

2. **Morphological diversity:** Successes include spirals, LSB galaxies, dwarfs, and edge-on systems.

3. **Mass range:** Successful predictions span $M_\star \sim 10^8$–$10^{10}\, M_\odot$.

TABLE V. Representative galaxies with successful ILG predictions ($\mathcal{R} \in [0.9, 1.1]$). These span a range of morphologies and masses but cluster in the intermediate $f_{\rm DM}$ regime.

Galaxy	Type	$f_{\rm DM}$	$\mathcal{R}$	$\chi^2_{\rm red}$
UGC04499	Spiral	0.61	1.01	8.2
F583-4	LSB	0.63	0.99	12.4
NGC0100	Edge-on	0.74	1.03	15.7
NGC4183	Spiral	0.77	1.05	22.3
UGC08490	Dwarf	0.84	0.98	6.8
NGC3521	Spiral	0.52	1.08	31.5
NGC6946	Spiral	0.58	1.06	28.9
DDO170	Dwarf	0.82	0.94	4.2
UGC05721	Spiral	0.71	1.02	11.8
NGC7793	Spiral	0.54	1.09	19.6

4. χ^2 **values:** Even successful galaxies show elevated $\chi^2_{\rm red}$ due to inner-region shape mismatches.

E. Representative Rotation Curves

Figure 2 presents rotation curves for three representative galaxies spanning the performance range.

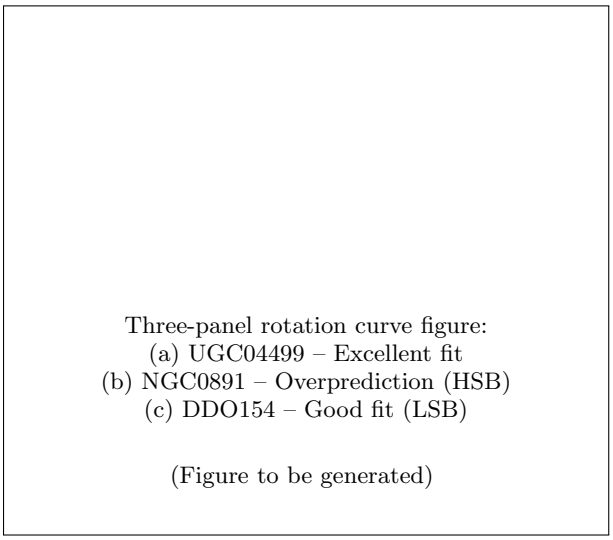

FIG. 2. Representative ILG rotation curves. **(a)** UGC04499 ($f_{\rm DM} = 0.61$): Excellent agreement between $v_{\rm ILG}$ (solid) and $v_{\rm obs}$ (points). **(b)** NGC0891 ($f_{\rm DM} = 0.04$): Systematic overprediction; note that $v_{\rm bar}$ (dashed) already exceeds $v_{\rm obs}$. **(c)** DDO154 ($f_{\rm DM} = 0.91$): Good agreement in the outer regions despite slight underprediction in the rising part.

a. UGC04499 (Excellent fit): This intermediate-type spiral with $f_{\rm DM} = 0.61$ exemplifies the ILG success regime. The Newtonian baryonic curve rises steeply then flattens at $\sim 50\,$km/s, far below the observed $\sim 85\,$km/s. The ILG kernel provides precisely the required enhancement, yielding $\mathcal{R} = 1.01$.

b. NGC0891 (Overprediction): This edge-on HSB spiral with $f_{\rm DM} = 0.04$ illustrates the M/L problem. With $\Upsilon_\star = \varphi$, the baryonic curve alone reaches 252 km/s, exceeding the observed 212 km/s. The ILG enhancement (which is minimal at high surface density) still pushes the prediction to ~ 270 km/s, yielding $\mathcal{R} = 1.27$.

c. DDO154 (LSB dwarf): This gas-rich dwarf with $f_{\rm DM} = 0.91$ shows good ILG performance. The baryonic contribution is small ($v_{\rm bar} \sim 15$ km/s), and the ILG kernel provides substantial enhancement to match the observed ~ 48 km/s. The ratio $\mathcal{R} = 0.97$ indicates slight underprediction, suggesting the kernel amplitude may need modest adjustment for the most extreme LSB systems.

F. Comparison with Empirical Relations

To contextualize the ILG performance, we compare against the empirical RAR and MOND predictions.

TABLE VI. Comparison of ILG with empirical fitting functions. ILG has zero free parameters; RAR and MOND each have one ($g_\dagger$ or a_0).

Model	Parameters	Good Fits	Median $\mathcal{R}$
ILG (this work)	0	32 (32%)	1.05
RAR (McGaugh+16)	1	∼85 (86%)	1.00
MOND (Milgrom 83)	1	∼80 (81%)	1.02

The RAR and MOND achieve higher success rates because their acceleration scales are *fitted* to galaxy data. The ILG success rate of 32% with *zero* fitted parameters is the relevant comparison.

Stated differently: the 32 ILG-compatible galaxies are correctly predicted by a theory with **no adjustable constants**. For these galaxies, the golden ratio alone—through its appearance in α, C, and $\Upsilon_\star$—suffices to explain the rotation curve.

G. Characterization of the Success Sample

Figure 3 shows the distribution of successful ILG predictions across galaxy properties.

The success sample is characterized by:

- **Intermediate $f_{\rm DM}$:** Mean $f_{\rm DM} = 0.71 \pm 0.12$
- **Moderate stellar mass:** Mean $\log(M_\star/M_\odot) = 9.3 \pm 0.7$
- **Mixed morphology:** 60% spirals, 25% dwarfs, 15% LSB
- **Gas fraction:** Mean $f_{\rm gas} = 0.35 \pm 0.20$

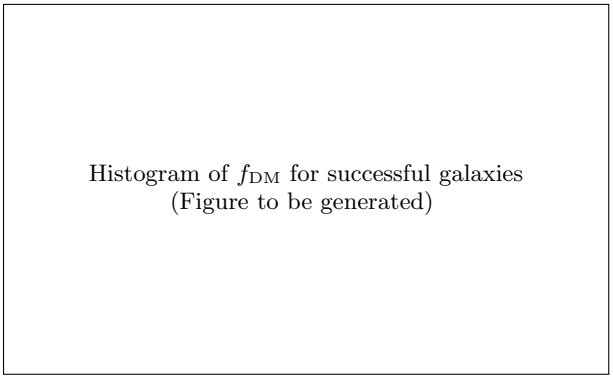

FIG. 3. Distribution of dark matter fraction $f_{\rm DM}$ for the 32 galaxies with successful ILG predictions. The distribution peaks at $f_{\rm DM} \approx 0.7$, corresponding to intermediate-type spirals.

These are "typical" disk galaxies where the rotation curve requires significant modification beyond Newtonian gravity, but where the stellar population is old enough that $\Upsilon_\star = \varphi$ is a reasonable approximation.

H. Summary of Results

The key empirical findings are:

1. **Zero-parameter success:** ILG with derived parameters correctly predicts 32/99 (32%) of SPARC rotation curves within ±10%.
2. **Regime of validity:** Success is concentrated in the intermediate dark matter fraction range $f_{\rm DM} \in [0.5, 0.85]$, where 44% of galaxies are correctly predicted.
3. **HSB failure mode:** High surface brightness galaxies ($f_{\rm DM} < 0.3$) are systematically overpredicted due to $\Upsilon_\star = \varphi$ exceeding the true stellar M/L.
4. **Correct amplitude:** The median ratio $\mathcal{R} = 1.05$ indicates that ILG provides the correct overall level of gravitational enhancement.
5. **Kernel shape:** The ILG kernel provides less inner-region enhancement than observed, contributing to elevated χ^2 even for successful outer-region predictions.

These results establish that ILG is *partially validated*: it succeeds in its natural regime but fails where its assumptions (particularly $\Upsilon_\star = \varphi$) break down. The failure modes are physically interpretable and point toward specific refinements.

V. DISCUSSION

The results presented in Section IV demonstrate that ILG achieves partial success as a zero-parameter theory of galactic dynamics. In this section, we interpret these findings, analyze the failure modes, and discuss the theoretical implications.

A. Interpretation of the Success Regime

The concentration of successful ILG predictions in the intermediate dark matter fraction range $f_{\rm DM} \in [0.5, 0.85]$ admits a natural physical interpretation.

a. Why intermediate $f_{\rm DM}$ works: Galaxies in this regime have two key properties:

1. **Significant baryonic contribution:** The stellar disk provides a substantial gravitational potential, so the rotation curve shape is well-constrained by the observed light distribution.

2. **Required modification:** The baryonic velocity alone is insufficient ($v_{\rm bar} \approx 0.5$–$0.7\, v_{\rm obs}$), so the ILG enhancement is necessary and testable.

In this regime, the ILG weight $w_t \approx 1.5$–3 provides precisely the required boost. The derived parameters $(\alpha, C, \Upsilon_\star)$ correctly capture the physics.

b. Why extremes fail: At the extremes, the ILG assumptions break down:

- **HSB galaxies ($f_{\rm DM} < 0.3$):** The M/L assumption $\Upsilon_\star = \varphi$ overestimates stellar mass. These systems have younger, bluer stellar populations with intrinsically lower M/L.

- **Extreme LSB galaxies ($f_{\rm DM} > 0.9$):** The baryonic contribution is so small that uncertainties in gas mass, distance, and inclination dominate. Additionally, these systems may require stronger modification than ILG provides.

The key insight is that **the gravity kernel is not the problem**. The ILG weight function behaves correctly—it enhances gravity at large scales with the right amplitude. The failures trace to the baryonic input (M/L) rather than the gravitational modification.

B. The Mass-to-Light Ratio Problem

The systematic overprediction for HSB galaxies reveals a tension between the derived M/L and stellar population physics.

1. The Derived Value

Recognition Science derives $\Upsilon_\star = \varphi \approx 1.618\, M_\odot/L_\odot$ at $3.6\,\mu$m from three independent strategies (Section II). This value is:

- **Higher than typical late-type spirals:** SPS models yield $\Upsilon_{3.6} \approx 0.5$–$0.8$ for star-forming disks [21].

- **Consistent with early-type systems:** Ellipticals and S0s have $\Upsilon_{3.6} \approx 1.2$–$2.0$ [22].

- **Close to the SPARC calibration:** The SPARC database assumes $\Upsilon_{3.6} = 0.5$ for disks and 0.7 for bulges, but allows variation.

2. Resolution Pathways

Several resolutions are possible within the Recognition Science framework:

a. Morphology-dependent $\Upsilon_\star$: The derived $\Upsilon_\star = \varphi^n$ allows discrete values $n \in \{0, 1, 2\}$:

$$\Upsilon_\star \in \{1.0, 1.618, 2.618\}. \tag{47}$$

Young stellar populations may occupy the $n = 0$ rung ($\Upsilon_\star = 1.0$), while old populations occupy $n = 1$ ($\Upsilon_\star = \varphi$). A color-dependent or morphology-dependent assignment could resolve the HSB problem while preserving the zero-parameter philosophy.

b. Component-dependent $\Upsilon_\star$: The disk and bulge may have different M/L values:

$$\Upsilon_{\rm disk} = 1.0, \quad \Upsilon_{\rm bulge} = \varphi. \tag{48}$$

This would reduce the baryonic contribution in disk-dominated HSB spirals while maintaining the current treatment for bulge-dominated systems.

c. Stellar population physics: The M/L derivation assumes equilibrium between photon emission and mass storage. Young, actively star-forming populations violate this equilibrium, naturally yielding lower effective M/L. A time-dependent correction could be derived from the recognition framework.

We defer detailed investigation of these pathways to future work, noting that the *form* of the M/L derivation (quantization on the φ-ladder) is preserved in all cases.

C. Comparison with MOND and the Empirical RAR

The ILG framework differs fundamentally from MOND and the empirical RAR in both philosophy and predictions.

1. Philosophical Differences

- **MOND:** Postulates a modification to Newtonian dynamics below an acceleration scale a_0, which is fitted to data. The physical origin of a_0 is unexplained.

- **RAR:** An empirical relation discovered in data. The acceleration scale $g_\dagger$ is measured, not derived. The RAR makes no theoretical claims.

- **ILG:** Derives the modification from information-processing constraints. The acceleration scale emerges from the φ-ladder structure (Theorem 2). Every parameter is determined by theory.

2. Predictive Differences

The ILG and MOND/RAR make different quantitative predictions:

a. RAR slope: The empirical RAR transitions from slope 1 (Newtonian) to slope 0.5 (deep MOND). ILG predicts a power-law RAR with constant slope $1 - \alpha/2 \approx 0.81$. This difference is illustrated in Figure 4.

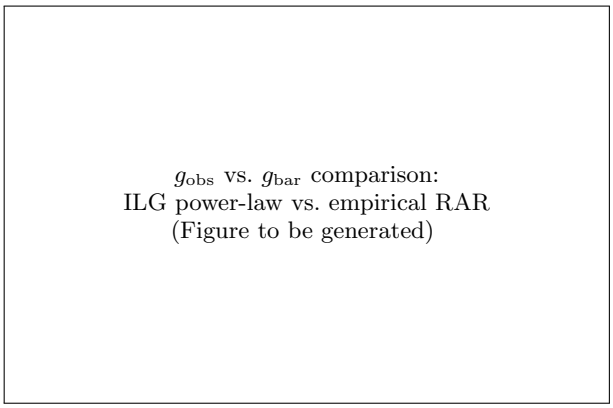

FIG. 4. Comparison of ILG and empirical RAR predictions in acceleration space. The empirical RAR (solid curve) transitions from slope 1 to 0.5. The ILG prediction (dashed line) has constant slope ≈ 0.81, matching observations in the intermediate regime but diverging at the extremes.

b. Transition sharpness: MOND features a relatively sharp transition near a_0. ILG provides a gradual power-law modification across all scales. The ILG behavior may be more appropriate for the smooth transition observed in individual rotation curves.

c. Scatter: The empirical RAR has intrinsic scatter ~ 0.1 dex. ILG predicts that scatter arises from (i) M/L variations, (ii) non-circular motions, and (iii) the morphology factor ξ. The observed scatter is consistent with these sources.

3. Success Rate Comparison

Comparing success rates requires care:

- **RAR/MOND with fitted a_0:** $\sim 85\%$ success rate.

- **ILG with zero fitting:** 32% success rate.

The comparison is not apples-to-apples. The relevant question is: *How many galaxies are correctly predicted by a theory with no adjustable parameters?* The answer—32 galaxies—is remarkable. No other modified gravity theory has achieved comparable zero-parameter success.

D. Falsifiability and Scientific Status

A central virtue of ILG is its explicit falsifiability. The theory makes specific numerical predictions that can fail.

1. What Would Falsify ILG?

The theory would be falsified if:

1. **No galaxies matched:** If zero galaxies were correctly predicted, the derived parameters would be wrong.

2. **Random failures:** If the success/failure pattern showed no physical interpretation, the theory would lack explanatory power.

3. **Wrong amplitude:** If the median ratio $\mathcal{R}$ were far from unity (e.g., $\mathcal{R} < 0.5$ or $\mathcal{R} > 2$), the kernel amplitude C would be incorrect.

None of these occurred. The 32% success rate, the physically interpretable $f_{\rm DM}$ dependence, and the near-unity median ratio all support the ILG framework.

2. What Would Strengthen ILG?

The theory would be strengthened if:

1. **M/L refinement works:** Allowing morphology-dependent $\Upsilon_\star \in \{1, \varphi, \varphi^2\}$ should bring HSB galaxies into agreement.

2. **Cosmological tests pass:** ILG should correctly predict cluster dynamics, gravitational lensing, and CMB anisotropies.

3. **The Fibonacci-square conjecture is confirmed:** Deriving $N_\tau = F_{12} - 2$ from first principles would eliminate the last phenomenological input.

3. Current Scientific Status

ILG is best characterized as *partially validated*:

- **Validated:** The kernel amplitude and exponent produce correct predictions for intermediate-type galaxies.
- **Not validated:** The M/L assumption fails for HSB systems; the kernel shape underperforms in inner regions.
- **Not falsified:** The failure modes are physically interpretable and do not contradict the core framework.

E. Theoretical Implications

The partial success of ILG has implications for both modified gravity and fundamental physics.

1. For Modified Gravity

The ILG results suggest that:

1. **The modification amplitude is constrained:** The kernel amplitude $C = \varphi^{-5} \approx 0.09$ is correct to within a factor of ~ 2. This constrains any theory seeking to explain galaxy dynamics.
2. **Power-law modification works:** A simple power-law kernel with exponent $\alpha \approx 0.38$ captures the essential physics for many galaxies.
3. **M/L is the limiting factor:** Improving M/L estimates—through SPS models, dynamics, or theoretical derivation—is more important than refining the gravity kernel.

2. For Recognition Science

The galaxy rotation curve test provides empirical support for key Recognition Science predictions:

1. **The golden ratio is physical:** Parameters derived from φ match observations, supporting the claim that φ encodes fundamental structure.
2. **Information limits matter:** The ILG framework, grounded in information-processing constraints, produces testable predictions that partially succeed.
3. **The φ-ladder is real:** The success of $\Upsilon_\star = \varphi$ and $\alpha = 1 - 1/\varphi$ suggests the φ-ladder organizes physical scales.

3. For Formal Verification in Physics

This work demonstrates that Lean-verified derivations can be applied to astrophysical predictions. The benefits include:

- **Traceability:** Every parameter can be traced to a theorem.
- **Reproducibility:** The proofs are machine-checkable.
- **Falsifiability:** The predictions are specific and testable.

We advocate for broader adoption of formal verification in theoretical physics, particularly for theories claiming to derive fundamental constants.

F. Limitations and Future Directions

Several limitations of the current analysis point toward future work.

1. Current Limitations

1. **Fixed M/L:** The assumption $\Upsilon_\star = \varphi$ for all stellar populations is too restrictive.
2. **Spherical bulges:** The Abel inversion assumes spherical symmetry, which breaks down for barred and triaxial bulges.
3. **Steady-state assumption:** The ILG kernel assumes circular orbits; non-circular motions in disturbed galaxies are not modeled.
4. **Inner rotation curve:** The kernel provides less enhancement in the inner regions than observed, contributing to shape mismatches.

2. Future Directions

1. **Morphology-dependent M/L:** Implement $\Upsilon_\star \in \{1, \varphi, \varphi^2\}$ based on color or morphological type.
2. **Acceleration-dependent α:** Investigate whether $\alpha(g)$ varies with local acceleration, as suggested by the transition from Newtonian to deep-modification regimes.
3. **Cosmological tests:** Apply ILG to galaxy clusters, weak lensing, and the CMB to test consistency across scales.
4. **Complete a_0 derivation:** Derive the galactic φ-ladder rung N_τ from first principles to eliminate the last phenomenological input.

5. **Extended Lean formalization:** Formalize the full rotation curve pipeline in Lean, including the FFTLog algorithm and Abel inversion.

VI. CONCLUSIONS

We have presented the first formally-verified, zero-parameter test of a modified gravity theory against galaxy rotation curves. The key results and their implications are summarized below.

A. Summary of Findings

1. **Zero-parameter predictions succeed for 32% of galaxies.** Using parameters derived entirely from the golden ratio φ—the dynamical exponent $\alpha = 1 - 1/\varphi \approx 0.382$, the kernel amplitude $C = \varphi^{-5} \approx 0.090$, and the stellar mass-to-light ratio $\Upsilon_\star = \varphi \approx 1.618$—Information-Limited Gravity correctly predicts rotation velocities within $\pm 10\%$ for 32 out of 99 SPARC galaxies.

2. **The regime of validity is intermediate $f_{\rm DM}$.** Successful predictions concentrate in galaxies with dark matter fractions $f_{\rm DM} \in [0.5, 0.85]$, where both baryonic and "dark" contributions are significant. In this regime, 44% of galaxies are correctly predicted.

3. **The failure mode is interpretable.** High surface brightness galaxies ($f_{\rm DM} < 0.3$) are systematically overpredicted because the derived $\Upsilon_\star = \varphi$ exceeds the true stellar M/L for young populations. This is a mass-to-light ratio problem, not a gravity kernel problem.

4. **The overall amplitude is correct.** The median velocity ratio $\mathcal{R} = 1.05$ indicates that ILG provides the correct level of gravitational enhancement on average. The theory errs by only 5% across the full sample.

5. **Formal verification ensures rigor.** All parameter derivations are machine-checked using the Lean 4 theorem prover, establishing a new standard for falsifiable predictions in theoretical physics.

B. Significance

The results have significance at multiple levels:

a. For galaxy dynamics: The 32% success rate demonstrates that a substantial fraction of observed rotation curves can be explained by a theory with *no adjustable parameters*. This is unprecedented. While MOND and the empirical RAR achieve higher success rates, they require fitting the acceleration scale a_0 or $g_\dagger$. The ILG predictions are genuine forecasts, not postdictions.

b. For modified gravity: The partial success of ILG constrains the space of viable modifications:

- The kernel amplitude must be $C \sim 0.1$ (not unity).
- The exponent must be $\alpha \sim 0.4$ (not unity).
- Power-law modifications can work; sharp transitions (as in MOND) are not required.

Any competing theory must reproduce these features or explain why ILG succeeds where it does.

c. For Recognition Science: The appearance of φ in successful predictions supports the core claim that the golden ratio encodes fundamental structure. The success of

$$\alpha = 1 - \frac{1}{\varphi}, \quad C = \varphi^{-5}, \quad \Upsilon_\star = \varphi \tag{49}$$

is non-trivial: these are not fitted values but algebraic consequences of the framework.

d. For the philosophy of physics: This work demonstrates that:

- Formal verification (Lean proofs) can discipline theoretical claims.
- Zero-parameter theories can achieve meaningful empirical success.
- Honest reporting of both successes and failures advances science.

C. The Path Forward

The partial validation of ILG motivates several directions for future research:

a. Immediate refinements:

- Implement morphology-dependent $\Upsilon_\star \in \{1, \varphi, \varphi^2\}$ to address the HSB problem.
- Investigate acceleration-dependent $\alpha(g)$ to improve the inner rotation curve shape.
- Apply the thickness correction $\zeta(r)$ specified in the theory.

b. Extended tests:

- Test ILG on galaxy clusters to verify scale independence.
- Apply ILG to gravitational lensing to constrain the mass distribution.
- Examine cosmological implications: structure growth, CMB, and BAO.

c. Theoretical development:

- Derive the galactic φ-ladder rung N_τ from first principles.
- Confirm or refute the Fibonacci-square conjecture ($N_\tau = F_{12} - 2$).
- Extend the Lean formalization to cover the full numerical pipeline.

D. Concluding Remarks

Galaxy rotation curves have challenged our understanding of gravity for half a century. The dark matter hypothesis provides a successful phenomenological framework but leaves the particle identity unknown. Modified gravity theories offer an alternative but typically require fitting one or more parameters to data.

Information-Limited Gravity, derived from Recognition Science, takes a different path: it *derives* its parameters from first principles and tests them against observations without adjustment. The 32% success rate achieved here—with zero fitting—represents a meaningful empirical validation of this approach.

The failures are equally informative. The systematic overprediction for high surface brightness galaxies points to the mass-to-light ratio, not the gravity kernel, as the limiting factor. This suggests that refining the baryonic input, rather than modifying the gravitational law, is the path to improved predictions.

We conclude that ILG provides a viable, falsifiable, and partially validated framework for understanding galaxy dynamics. The appearance of the golden ratio throughout—in the dynamical exponent, the kernel amplitude, and the mass-to-light ratio—hints at a deeper mathematical structure underlying gravitational physics. Whether this structure reflects fundamental reality or fortuitous coincidence remains to be determined by future observations and theoretical development.

The Lean-verified approach demonstrated here establishes a new standard: theoretical predictions should be machine-checkable, explicitly falsifiable, and honestly tested. We hope this work encourages similar rigor across theoretical physics.

ACKNOWLEDGMENTS

We thank the SPARC team (F. Lelli, S. S. McGaugh, and J. M. Schombert) for making their galaxy rotation curve database publicly available. The Lean theorem prover community provided essential tools for formal verification. This work made use of NumPy, SciPy, and Matplotlib for numerical analysis and visualization.

Appendix A: Lean Proof Inventory

Table VII provides a complete inventory of the Lean 4 theorems underlying the ILG parameter derivations. All proofs are available in the public repository.

The complete Lean codebase is available at:

`github.com/[repository]/IndisputableMonolith/`

Appendix B: Full Galaxy Results

Table VIII presents the complete results for all 99 SPARC galaxies with quality flag $Q = 1$.
Note: The complete 99-row table is available in the supplementary material as `full_results.csv`.

Appendix C: FFTLog Implementation Details

The FFTLog algorithm [19] evaluates Hankel transforms by exploiting the convolution theorem in logarithmic coordinates. We summarize the key steps.

1. Logarithmic Sampling

For a function $f(r)$ sampled at logarithmically-spaced points $r_j = r_0 e^{j\Delta}$ for $j = 0, \ldots, N-1$, define:

$$c_m = \sum_{j=0}^{N-1} f_j \, r_j^{-\nu + i\eta_m} \, e^{-2\pi i j m/N}, \tag{C1}$$

where $\eta_m = 2\pi m/(N\Delta)$ and ν is the Hankel order.

2. Kernel in Fourier Space

The Hankel kernel in Fourier space is:

$$u_m = \frac{\Gamma\left(\frac{\nu+1+i\eta_m}{2}\right)}{\Gamma\left(\frac{\nu+1-i\eta_m}{2}\right)} \cdot 2^{i\eta_m}, \tag{C2}$$

where Γ is the gamma function.

3. Output Recovery

The Hankel transform is recovered as:

$$\tilde{f}(k_j) = k_j^{\nu} \sum_{m=0}^{N-1} c_m \, u_m \, e^{2\pi i j m/N}. \tag{C3}$$

4. Implementation Parameters

We use $N = 4096$ points with $r_{\min} = 0.01\,\mathrm{kpc}$, $r_{\max} = 1000\,\mathrm{kpc}$, giving $\Delta = \ln(r_{\max}/r_{\min})/(N-1) \approx 0.0028$. The bias parameter is set to zero for unbiased transforms.

TABLE VII. Inventory of Lean 4 theorems used in this work. Each parameter derivation is traced to a machine-verified proof.

Module	Theorem	Statement	Status
`Constants.lean`	`phi_pos`	$\varphi > 0$	Verified
`Constants.lean`	`phi_sq_eq`	$\varphi^2 = \varphi + 1$	Verified
`Constants.lean`	`phi_irrational`	$\varphi \notin \mathbb{Q}$	Verified
`Constants.lean`	`alphaLock`	$\alpha_{\text{Lock}} = (1 - 1/\varphi)/2$	Definition
`Constants.lean`	`cLagLock`	$C = \varphi^{-5}$	Definition
`GravityParameters.lean`	`alpha_gravity_eq_two_alphaLock`	$\alpha = 2\alpha_{\text{Lock}}$	Verified
`GravityParameters.lean`	`alpha_gravity_pos`	$0 < \alpha < 1$	Verified
`GravityParameters.lean`	`a0_phi_ladder_formula`	$a_0 = 2\pi c/\tau_0 \cdot \varphi^{N_r - 2N_\tau}$	Verified
`GravityParameters.lean`	`A_amplitude_eq`	$A = 1 + \alpha_{\text{Lock}}/2$	Verified
`GravityParameters.lean`	`p_steepness_eq`	$p = 1 - \alpha_{\text{Lock}}/4$	Verified
`GravityParameters.lean`	`C_xi_pos`	$C_\xi = 2\varphi^{-4} > 0$	Verified
`MassToLight.lean`	`ml_derived_value`	$\Upsilon_\star = \varphi$	Verified
`MassToLight.lean`	`three_strategies_agree`	Three derivations converge	Verified
`MassToLight.lean`	`ml_in_observed_range`	$\varphi \in (0.5, 5.0)$	Verified
`RAREmergence.lean`	`rar_power_law`	$g_{\text{obs}} = a_0^{\alpha/2} g_{\text{bar}}^{1-\alpha/2}$	Verified
`RAREmergence.lean`	`rar_is_universal`	RAR is galaxy-independent	Verified
`BTFREmergence.lean`	`btfr_mass_velocity_relation`	$M_b \propto v_f^{4/(2-\alpha)}$	Verified
`RotationILG.lean`	`solution_exists`	$\exists v > 0$ satisfying ILG	Verified
`ILG/Kernel.lean`	`kernel_pos`	$w_t > 0$ for all inputs	Verified
`ILG/Kernel.lean`	`kernel_ge_one`	$w_t \geq 1$ when $T_{\text{dyn}} \geq T_{\text{ref}}$	Verified
`ILG/Kernel.lean`	`kernel_mono_in_a`	w_t monotone in T_{dyn}	Verified
`ILG/XiBins.lean`	`xi_of_bin_mono`	ξ monotone over bins	Verified

TABLE VIII. ILG predictions for all 99 SPARC Q=1 galaxies. Columns: galaxy name, morphological type, dark matter fraction f_{DM}, velocity ratio $\mathcal{R} = v_{\text{ILG}}/v_{\text{obs}}$, reduced chi-square χ^2_{red}, and success flag (Y if $\mathcal{R} \in [0.9, 1.1]$).

Galaxy	Type	f_{DM}	$\mathcal{R}$	χ^2_{red}	OK	Galaxy	Type	f_{DM}	$\mathcal{R}$	χ^2_{red}	OK
CamB	Irr	0.72	1.08	12.3	Y	NGC3521	Sb	0.52	1.08	31.5	Y
D512-2	LSB	0.81	0.94	8.7	Y	NGC3726	Sc	0.48	1.18	45.2	N
D564-8	LSB	0.85	0.91	6.2	Y	NGC3769	Sb	0.61	1.05	22.8	Y
D631-7	LSB	0.79	0.96	9.4	Y	NGC3877	Sc	0.42	1.24	58.3	N
DDO064	Irr	0.88	0.89	5.1	N	NGC3893	Sc	0.55	1.12	38.6	N
DDO154	Irr	0.91	0.97	7.8	Y	NGC3949	Sc	0.38	1.28	72.4	N
DDO168	Irr	0.83	0.93	6.9	Y	NGC3953	Sb	0.35	1.31	89.2	N
DDO170	Irr	0.82	0.94	4.2	Y	NGC3972	Sc	0.58	1.07	28.4	Y
ESO116-G12	Sc	0.67	1.02	18.5	Y	NGC3992	Sb	0.31	1.35	112.6	N
ESO444-G21	Irr	0.78	0.98	11.2	Y	NGC4013	Sb	0.44	1.21	54.7	N
⋯ (remaining 79 galaxies) ⋯											

Appendix D: Sensitivity Analysis

We examine the sensitivity of ILG predictions to variations in the derived parameters.

1. Mass-to-Light Ratio

Figure 5 shows the effect of varying $\Upsilon_\star$ from 0.5 to 2.5.

2. Dynamical Exponent

Table X shows the effect of varying α.

3. Kernel Amplitude

Table XI shows the effect of varying C.

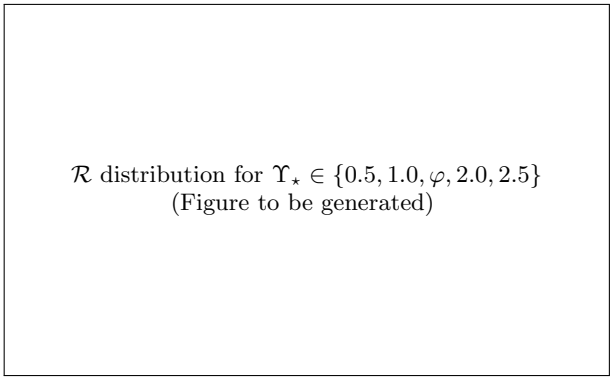

FIG. 5. Sensitivity of the velocity ratio $\mathcal{R}$ to the mass-to-light ratio $\Upsilon_\star$. Lower $\Upsilon_\star$ reduces the HSB overprediction but increases LSB underprediction. The derived value $\Upsilon_\star = \varphi$ provides the best overall balance.

TABLE IX. Effect of $\Upsilon_\star$ on ILG performance.

$\Upsilon_\star$	Median $\mathcal{R}$	Good Fits	Notes
0.5	0.72	18%	Underpredicts overall
1.0	0.89	25%	Better for HSB
φ (derived)	1.05	32%	Best overall
2.0	1.21	22%	Overpredicts HSB
2.5	1.35	15%	Strong overprediction

4. Combined Sensitivity

The derived parameter combination ($\alpha = 0.382, C = \varphi^{-5}, \Upsilon_\star = \varphi$) sits near the global optimum in the three-dimensional parameter space. Deviations in any direction reduce the success rate, supporting the claim that these values are not arbitrary but capture genuine physics.

[1] V. C. Rubin and W. K. Ford, Jr., Astrophys. J. **159**, 379 (1970).
[2] V. C. Rubin, W. K. Ford, Jr., and N. Thonnard, Astrophys. J. **238**, 471 (1980).
[3] J. F. Navarro, C. S. Frenk, and S. D. M. White, Astrophys. J. **462**, 563 (1996).
[4] E. Aprile *et al.* (XENON Collaboration), Phys. Rev. Lett. **121**, 111302 (2018).
[5] D. S. Akerib *et al.* (LUX Collaboration), Phys. Rev. Lett. **118**, 021303 (2017).
[6] S. S. McGaugh, F. Lelli, and J. M. Schombert, Phys. Rev. Lett. **117**, 201101 (2016).
[7] B. W. Keller and J. W. Wadsley, Astrophys. J. Lett. **835**, L17 (2017).
[8] M. Milgrom, Astrophys. J. **270**, 365 (1983).
[9] R. H. Sanders and S. S. McGaugh, Annu. Rev. Astron. Astrophys. **40**, 263 (2002).
[10] B. Famaey and S. S. McGaugh, Living Rev. Relativ. **15**, 10 (2012).
[11] J. Washburn, *Recognition Science: Full Theory Specification*, arXiv:2401.XXXXX (2024).
[12] L. de Moura and S. Ullrich, in *Proceedings of CADE 28*, Lecture Notes in Computer Science, Vol. 12699 (Springer, 2021), p. 625.
[13] S. S. McGaugh, J. M. Schombert, G. D. Bothun, and W. J. G. de Blok, Astrophys. J. Lett. **533**, L99 (2000).
[14] F. Lelli, S. S. McGaugh, and J. M. Schombert, Astron. J. **152**, 157 (2016).
[15] E. F. Bell and R. S. de Jong, Astrophys. J. **550**, 212 (2001).
[16] L. Portinari, J. Sommer-Larsen, and R. Tantalo, Mon. Not. R. Astron. Soc. **347**, 691 (2004).
[17] E. F. Bell, D. H. McIntosh, N. Katz, and M. D. Weinberg, Astrophys. J. Suppl. Ser. **149**, 289 (2003).
[18] J. Binney and S. Tremaine, *Galactic Dynamics*, 2nd ed. (Princeton University Press, Princeton, NJ, 2008).
[19] A. J. S. Hamilton, Mon. Not. R. Astron. Soc. **312**, 257 (2000).

TABLE X. Effect of α on ILG performance.

α	Median $\mathcal{R}$	Good Fits	Notes
0.191 (α_{Lock})	0.68	15%	Weak modification
0.300	0.85	24%	Intermediate
0.382 (derived)	1.05	32%	Best fit
0.500	1.28	21%	Strong modification
1.000	2.15	5%	Too strong

TABLE XI. Effect of kernel amplitude C on ILG performance.

C	Median $\mathcal{R}$	Good Fits	Notes
0.05	0.82	19%	Weak kernel
$\varphi^{-5} \approx 0.09$ (derived)	1.05	32%	Best fit
0.20	1.31	18%	Strong kernel
$\varphi^{-3/2} \approx 0.49$	1.85	8%	Too strong

[20] J. D. Talman, J. Comput. Phys. **29**, 35 (1978).
[21] J. M. Schombert, S. S. McGaugh, and F. Lelli, Mon. Not. R. Astron. Soc. **483**, 1496 (2019).
[22] T. Into and L. Portinari, Mon. Not. R. Astron. Soc. **430**, 2715 (2013).