The Corona in the Well: A Gravitational Correction to the Quasar Hubble Diagram

Martin Scholl — Independent Researcher

It Is All One — Notes. July 2026 (readable edition)

Abstract. An idea this series once buried — that part of a source's redshift might be gravitational, born in a deep well rather than accumulated on the way — resurrects at the one rung where the physics genuinely supports it. Quasar cosmology is built on X-ray light that is born a few gravitational radii from a supermassive black hole: genuinely deep in a well. This note computes what the well does to that light (two effects, both from this series' own toolbox), shows that the standard quasar-distance method *amplifies* the X-ray effect two-and-a-half-fold, and finds that a modest, astrophysically plausible drift of coronal depth across the quasar population produces a bias of the same size, sign and location as the celebrated deviation of quasar cosmology from the standard model at high redshift — the anomaly behind "evolving dark energy" headlines. If so, that anomaly is general relativity in the engine, not new physics in the expansion. Tests are stated; every symbol is introduced before use.

1. The Idea, Reborn at the Right Rung

The local-well correction died for supernovae for cause: a white dwarf's explosion shines from millions of kilometres of ejected gas, where its own gravity reddens light by only a few parts in ten million. Nothing. But the idea was right and the object was wrong. A quasar is a supermassive black hole feeding: its ultraviolet light comes from the surrounding disk of hot matter, tens to hundreds of gravitational radii out — but its X-rays come from the *corona*, a compact region of energized plasma sitting at roughly five to ten Schwarzschild radii (the Schwarzschild radius $r_s = 2GM/c^2$ is the size the central mass would have as a bare black hole). Five Schwarzschild radii is genuinely deep: and modern cosmology uses exactly this light, because the relation between quasars' X-ray and ultraviolet brightness (Risaliti and Lusso) turns them into standardizable candles reaching seven times deeper in redshift than any supernova.

2. What the Well Does to the Light

Two effects, both already in this series' toolbox. First, *gravitational redshift dimming*: light climbing from radius r loses energy, and arrives at a rate slowed by the same factor — together dimming the flux by $(1 - r_s/r)$. Second, *the exit cone* (Paper 1, Section 6.1): close to a black hole, not all directions lead out; a fraction of the photons emitted isotropically simply falls in. Combining the two, per emission radius:

emission radius	gravitational redshift	photons captured	total dimming (magnitudes)
3 r_s	0.225	14.6%	0.612
5 r_s	0.118	5.7%	0.306
10 r_s	0.054	1.5%	0.131
25 r_s	0.021	0.3%	0.047
100 r_s	0.005	0.0%	0.011

A corona at five Schwarzschild radii loses **0.31 magnitudes** to its own well; the ultraviolet disk at twenty-five loses 0.05. These are not small numbers, and they are not speculative physics — they are the standard picture of the quasar engine, priced with Schwarzschild optics.

3. The Amplifier

The Risaliti–Lusso method infers a quasar's distance by comparing its X-ray flux against its ultraviolet flux through their measured statistical relation. Propagating a dimming through that construction (differentiating the distance formula the method uses), one finds:

bias in inferred distance modulus = 2.5 × (X-ray dimming) − 1.5 × (UV dimming)

The construction amplifies X-ray dimming two-and-a-half-fold. For the typical geometry above, the standing bias is +0.70 magnitudes — but a bias *constant across all quasars* is silently absorbed when the relation is calibrated, exactly the lesson of the buried supernova correction. What cosmology feels is only the *drift* of coronal depth across the sample.

4. The Scenario That Matches the Anomaly

The quasar Hubble diagram famously deviates from the standard model at redshifts above about 3, by 0.3–0.6 magnitudes in the direction of *dimmer and farther* — the claim behind "evolving dark energy." Now price a drift: if the effective coronal radius across the sample slides from ten Schwarzschild radii at low redshift to five at high redshift, the induced bias is **+0.44 magnitudes of apparent dimming at high redshift**; ten-to-four gives +0.70. Direction and magnitude both match the claimed anomaly. And a selection mechanism lies ready to hand: surveys at high redshift see only the most luminous engines, and the most luminous, fastest-feeding engines are precisely those whose coronae are observed (by X-ray timing and microlensing) to be most compact. No cosmic conspiracy required — the telescope's own selection walks the sample down the well.

The hypothesis, plainly labeled: the high-redshift quasar anomaly may be gravitational redshift plus exit-cone capture of coronal X-rays, amplified by the distance construction and driven by luminosity selection — relativity in the engine, not new physics in the expansion. If so, it is a systematic to be modeled before the quasar record can referee *any* cosmology: the standard model and metric D alike (whose own signature in that band, −0.15 to +0.07 magnitudes, is smaller than this effect).

5. Caveats, Honestly

The corona is not a static emitter — motions and beaming modify the static-limit numbers; the table is the leading-order skeleton. The ultraviolet disk is itself mildly relativistic. The amplification factor depends on the measured slope of the X-ray/UV relation. And whether this specific bias has been proposed and tested in the quasar-cosmology literature must be checked in the literature pass this corpus owes — the idea is simple enough that someone may own it.

6. Tests

(1) Correlate the quasar Hubble-diagram residuals with independent measures of coronal compactness (X-ray variability timescales, the X-ray-to-UV spectral index, feeding rate): the hypothesis predicts the residual is an astrophysics axis, not a cosmology axis. (2) Re-fit the quasar record with a coronal-depth term included: the high-redshift anomaly should shrink toward the standard model — or toward metric D; the fit can carry both. (3) Gravitationally microlensed quasars yield direct coronal sizes: the subsample with measured coronae is the calibration set. Each test uses data that already exists.

7. Bottom Line

The redshift correction buried at the supernova rung resurrects, correctly, at the quasar rung — with this series' own exit-cone formula as its engine and a ×2.5 amplifier waiting inside the standard method. At the supernova rung the well was a millionth of a magnitude; at the quasar rung it is a third of a magnitude before amplification — and it sits exactly where modern cosmology reports its most publicized anomaly.

References

G. Risaliti and E. Lusso, Nature Astronomy 3, 272 (2019); E. Lusso et al., A&A 642, A150 (2020); quasar coronal sizes from X-ray reverberation and microlensing literature; and the papers and notes of this series (Paper 1 §6.1 — the exit cone; the cosmology paper; the Ledger of the Way). *(Citations from memory; the literature-verification pass applies.)*