Counteracting the Camera's Exaggeration of Lens Darkness: Perceptual Display Compensation for Through-Lens Photography

Ken Steiner
Banana Research, Banana Sport Optics, Bäch, Switzerland
bananaoptics.com/blogs/research

27 July 2026 · Technical note II, revision 2

Abstract. A camera at locked exposure records the physical truth about a tinted lens: a 15 % transmission lens yields a photograph at 15 % brightness. A human wearing the same lens experiences nothing of the sort, because the visual system adapts within moments, reopening the pupil and rescaling retinal response. Presented naively, through-lens photography therefore *exaggerates* darkness and misinforms exactly the comparison it is meant to serve. We derive a family of display gains applied in linear light, parameterised by a single exponent on measured relative transmission, that spans the physically honest rendering, an eye-adapted rendering based on a square-root brightness law, and a fully luminance-equalised rendering that isolates tint. We state the model's assumptions and propose a psychophysical matching experiment to calibrate the adaptation exponent against real lenses. The camera's companion exaggeration, the colour cast, which vision also partly discounts, is analysed and deliberately left uncompensated: tint reshaping is the lens's function and survives chromatic adaptation.

1 The problem

All through-lens frames in this study were captured at one locked exposure per scene, set on the bare sensor (companion note, *Camera-Relative Transmission Measurement*). Locking is what makes frames comparable, but it also means the photograph inherits the lens's full attenuation: the darker the lens, the darker the file, linearly. The eye behaves differently. Within seconds of putting on a dark lens, pupil dilation and retinal gain control recover much of the apparent brightness of the scene; what the wearer keeps noticing is the tint and the contrast rendering, not a world at 15 % brightness. A publication that shows locked-exposure frames without comment therefore systematically overstates how dark every lens feels in use, and overstates it most for the darkest lenses, precisely the ones a buyer is most cautious about.

2 Display gains in linear light

Let I be the linearly developed camera image of a scene through a lens with measured relative transmission $\mathrm{relT} \in (0, 1]$, and let S be a global white point chosen from the brightest frame of the scene so that no frame clips. The displayed image is

$$D = f_{\mathrm{sRGB}}(S \cdot \mathrm{relT}^{\gamma} \cdot I), \qquad (1)$$

where f_{sRGB} is the standard encoding transfer function [4] and the exponent γ selects the rendering:

Mode	γ	Displayed brightness
True exposure	0	$\propto$ relT (untouched)
Eye adapted	$-\frac{1}{2}$	$\propto \mathrm{relT}^{1/2}$
Colour only	-1	equal for all lenses

Because the gain acts on linear data before encoding, chromaticity is preserved: the modes differ in luminance placement only. Figure 1 shows one frame under all three renderings.

2.1 True exposure ($\gamma = 0$)

The untouched frame is the physical reference: what a light meter reports, and the rendering under which the measurement of the companion note is defined. It is deliberately retained as a mode because honesty about the baseline is what makes the compensated modes interpretable.

2.2 Eye adapted ($\gamma = -1/2$)

Steady-state brightness perception compresses luminance approximately as a power law; Stevens-type exponents for brightness cluster between roughly 0.33 and 0.5 depending on adaptation state and surround. Choosing the conservative end, perceived brightness scales as $L^{1/2}$, so a lens that attenuates luminance by relT feels darker by only $\mathrm{relT}^{1/2}$. Displaying $\mathrm{relT}^{-1/2} \cdot I$ places each lens at that perceived level: a 16 % lens is shown at 40 % brightness, not 16 %. The choice of the conservative exponent is deliberate on a second ground: it undercompensates rather than overcompensates, so dark lenses still read as darker, and the rendering never claims a dark lens is effectively free of cost. What no scalar gain can restore, and what therefore remains visible in this mode, is the part adaptation cannot undo: the spectral tint and the contrast rendering.

2.3 Colour only ($\gamma = -1$)

Full luminance equalisation removes darkness from the comparison entirely. This is not a claim about perception; it is an analysis tool that isolates the one variable the other modes leave entangled, the colour cast, for tint-to-tint comparison of what each lens does to greens, browns and skies.

3 Worked example

Prizm Trail measured relT = 0.44 in the mixed-light forest scene. Its three renderings apply gains of 1, $0.44^{-1/2} = 1.51$ and $0.44^{-1} = 2.27$ in linear light, displaying the frame at 44 %, 66 % and 100 % of the scene's reference brightness respectively. An implementation error in an early build applied $\gamma = 1$ instead of $\gamma = 0$ for true exposure, darkening every frame by its transmission a second time (relT^2 overall); the published mode ratios make such errors externally checkable, since the luminance ratio between true exposure and eye adapted must equal $\mathrm{relT}^{1/2}$ for every frame.

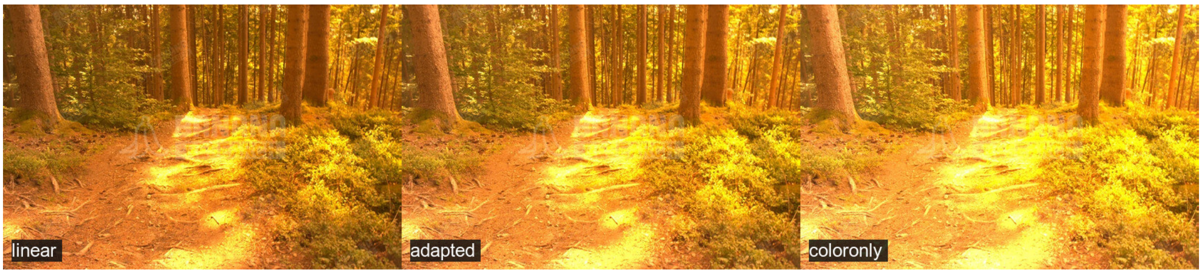

Figure 1: One field frame (Prizm Trail, mixed-light forest scene) under the three renderings. Left: true exposure, the locked-exposure file as captured. Centre: eye adapted, brightened by $\mathrm{relT}^{-1/2}$. Right: colour only, luminance equalised so only the tint remains.

4 Colour: the exaggeration we do not correct

The camera exaggerates a second thing besides darkness. With white balance fixed, the photograph shows a lens's full colour cast indefinitely, while a wearer experiences chromatic adaptation: the visual system rescales its colour channels toward the prevailing illumination (von Kries adaptation [3]), and a moderate uniform cast largely disappears from awareness within minutes. By the logic of Section 2 one might therefore expect a colour-adapted rendering alongside the brightness-adapted one. We deliberately do not provide it, for two reasons.

First, adaptation to strongly chromatic, illuminant-like shifts is incomplete, and its degree varies with tint strength, duration and observer; without measuring it, any partial correction factor would be an invented number, and unlike the brightness case there is no conservative default that fails safe. Second, and more fundamentally, sport lens tints are not uniform casts. They are engineered band filters that reshape object colours *relative to one another*; that reshaping is the lens's marketed function. Von Kries adaptation renormalises the channel means but preserves the relative restructuring, so the perceptual effect a buyer is evaluating survives adaptation even as the average cast fades. Displaying the full cast therefore keeps the between-lens comparison faithful, at the acknowledged cost of overstating the absolute cast a wearer settles into.

A colour-adapted rendering is nevertheless implementable in the same linear pipeline as Eq. (1), replacing the scalar gain with a diagonal per-channel gain toward the scene neutral, scaled by an empirical adaptation degree. We defer it until that degree is measured; the natural instrument is an achromatic-adjustment experiment through the physical lenses, the chromatic sibling of the brightness matching proposed below.

5 Assumptions and proposed calibration

The model asserts a single global exponent for all lenses, scenes and observers, ignores the time course of adaptation (it models the settled state, not the first seconds after putting the lens on), ignores pupil-versus-retinal contributions individually, and ignores any tint-dependent brightness effects (a strongly chromatic field can appear brighter than an equiluminant neutral one, the Helmholtz–Kohlrausch effect [3], which would make warm lenses feel slightly brighter than the model predicts). These are acceptable for a comparison publication; they are stated so readers do not mistake the eye-adapted mode for a measurement.

The exponent is calibratable with modest equipment, and we propose this as the natural next measurement: an observer wears the physical lens and views the real scene, then adjusts the displayed brightness of the through-lens photograph on a screen until the two match; repeating across the specimen set and several observers yields an empirical exponent per lens directly comparable to the assumed $1/2$, and would reveal any tint dependence. A luminance-calibrated display or a spot luminance meter would upgrade the match from subjective equality to measured equality.

6 Data and code availability

Measured transmissions and the rendering pipeline (linear development, gain application, white-point selection) are part of the study's build system; the dataset is published as JSON alongside the study at bananaoptics.com/blogs/research. The measurement protocol behind relT is described in the companion note.

References

[1] S. S. Stevens, "On the psychophysical law," *Psychological Review* 64(3):153–181, 1957.

[2] J. C. Stevens and S. S. Stevens, "Brightness function: effects of adaptation," *Journal of the Optical Society of America* 53(3):375–385, 1963.

[3] M. D. Fairchild, *Color Appearance Models*, 3rd ed., Wiley, 2013. Chromatic adaptation and the Helmholtz–Kohlrausch effect.

[4] IEC 61966-2-1:1999, *Multimedia systems and equipment, colour measurement and management, Part 2-1: Colour management, default RGB colour space, sRGB*.