Episode Book

EPISODE #5725

Why Your Smart Light's Lowest Setting Is a Lie

PUBLISHED	4 October 2026 at 19:43 UTC
DURATION	20:40
LISTEN	myweirdprompts.com/episode/led-dimming-range-spec-sheet/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.2
BOOK GENERATED	19 October 2026 at 17:39 UTC

SYNOPSIS

Your smart bulb can hit a perfect 2700K and still snap off like a switch at 10% brightness. Dimming range is the parameter that never makes it onto the box, and it's the one that decides how a light actually feels to live with. We dig into what "stepless" really means, why 8-bit controllers turn the bottom of the dial into a staircase, and why pulse-width modulation flicker gets worse exactly where you want the light gentlest. Then we connect it to the circadian research: intensity may matter more than blue light content, which makes deep dimming the most underrated health feature in lighting.

TRANSCRIPT



Corn

Here's a thing I've noticed about our little series on lighting. We did color rendering for inspection work. We did Kelvin for circadian alignment. Both of those are printed right on the box, front and center, in nice big numbers with a little icon next to them.



Herman

And then there's the one nobody puts on the box.



Corn

Right. Daniel sent over something he's been chewing on, and it's about dimming range, which he calls the parameter that doesn't get much attention but is maybe the most important one for how a light actually feels to live with.



Herman

He's not wrong about the attention gap.



Corn

So the short version of what he's asking. He wants to go deep on dimming range as a purchasing parameter. He's noticed a weird split in the market: expensive lights with perfect Kelvin adjustability and smart integrations that only give you three brightness steps, and cheap lights that advertise "stepless" dimming, which he points out can mean very different things depending on how good the controller is. He wants to know what dimming range actually is, what numbers show up on spec sheets, and why granular control at low brightness is a genuine engineering problem rooted in LED technology. Then he ties it to the circadian research, where he says the emerging case is that degree of light exposure probably matters more than blue light content. And he flags a cohort he's part of, which I suspect is bigger than he thinks, of people who find the minimum brightness on their phone still too bright. So: what's making finer dimming easier for manufacturers, in smart lights and dumb ones?



Herman

That's four questions and a confession.



Corn

It's a Daniel prompt. Let's start with the spec sheet, because I think the term itself is doing less work than people assume.



Herman

Dimming range is the span between a fixture's maximum output and its minimum controllable output. That's it. The measure is how far down it can go before it stops dimming and just turns off.



Corn

And it's expressed two ways.



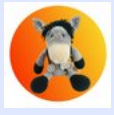
Herman

As a percentage of full brightness, or as a ratio. One percent dimming is the same thing as a hundred to one range. Point one percent is a thousand to one. Point zero one percent is ten thousand to one.



Corn

So the number on the box is really a floor. How low can it go before it quits.



Herman

And the tiers are pretty consistent across the industry, though you should treat these as conventions rather than gospel. Budget fixtures and a lot of smart bulbs sit around ten percent, which is a ten to one range. That's the spec where the light doesn't fade out, it snaps off.



Corn

You notice that immediately in a bedroom.



Herman

Mainstream dimmable LED drivers usually quote one percent. High end architectural and professional drivers claim point one. And specialist stuff goes to point zero one, ten thousand to one.



Corn

Which sounds incredible until you ask how it's measured.



Herman

That's the catch, and it's the thing Daniel is circling. The one percent on the sheet is often measured at a nominal condition. It tells you the driver can technically reach that point. It does not tell you the light behaves well when it gets there.



Corn

Two things the spec sheet leaves out, then. The curve and the resolution.



Herman

The curve is the shape of the response between minimum and maximum. Linear means the output tracks the input proportionally. Logarithmic, or square law, means the bottom of the range gets more of the control band.



Corn

And your eye wants the second one, because perception of brightness is roughly logarithmic.



Herman

A linear curve feels like nothing happens across the top ninety percent, and then everything happens in the last ten percent of the dial. Which is why cheap dimmers feel broken even when they're working as designed.



Corn

And resolution is where the "stepless" thing falls apart.



Herman

This is the point I want to land before we go anywhere else. An eight bit controller gives you two hundred and fifty six steps across the entire range. The whole span, top to bottom.



Corn

Split that across a hundred to one range and the bottom one percent of brightness gets two or three usable steps.



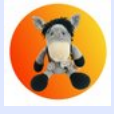
Herman

So "stepless" is a word about the interface, not about the light. The slider moves continuously. The output does not.



Corn

So Daniel's split market makes total sense. A luxury fixture can have gorgeous color control and a controller so coarse the bottom of the range is a staircase.



Herman

Stepless does not mean granular. Those are two different claims and the industry lets them blur.



Corn

So that's the spec sheet. Now the part I actually want to get into, which is why the bottom of the range is so hard to build.



Herman

There are two ways to dim a light emitting diode. Analog and pulse width modulation.



Corn

Analog first.



Herman

Analog dimming reduces the drive current to the emitter. You're literally asking the diode for less light, and it gives you less light. It's the intuitive approach.



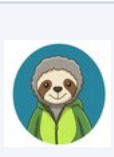
Corn

So why doesn't everything do it?



Herman

Because an LED doesn't behave like a resistor. It needs current limiting, and current is not something you set and forget. It's a dependent variable. It shifts with temperature, it shifts with the forward voltage of the junction, it shifts across the batch.



Corn

So you need a feedback loop.



Herman

You need to measure what's actually flowing and correct for it continuously. That's more components, more cost, more places for it to drift out of true. And there's the efficiency problem.



Corn

Which matters enormously on a battery.



Herman

Analog dimming wastes power in the driver itself. If you're running off a phone battery you cannot afford that. So the industry reaches for pulse width modulation, and pulse width modulation is where all the trouble starts.



Corn

Explain it.



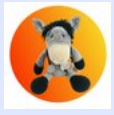
Herman

You don't reduce the current at all. You switch the emitter fully on and fully off, very fast, and you vary the proportion of time it spends on. Fifty percent brightness means it's on half the time.



Corn

And if you do that fast enough, the eye integrates it into a steady glow.



Herman

That's the theory. In practice you get flicker, and the flicker gets worse the deeper you dim. Because at low duty cycles the off periods are longer, and the longer the off period, the more likely your visual system notices it's there.



Corn

So the bottom of the range is exactly where the technique is weakest.



Herman

Raymond Soneira, who ran DisplayMate, put it in a way I've never been able to improve on. He said the flicker is often subliminal. People just feel unusually tired or uncomfortable and don't know why.



Corn

That's a bad sentence to read before bed.



Herman

There's a frequency dimension too. Pulse width modulation typically runs at four times the refresh rate. A hundred and twenty hertz panel means four hundred and eighty hertz flicker. The iPhone 12 was a sixty hertz display, so roughly two hundred and forty hertz.



Corn

Which is not that fast.



Herman

It's fast enough that most people can't consciously see it, and not fast enough that nobody reacts to it. Some people can sense even that.



Corn

The iPhone 12 Mini case is worth naming here because it's the closest thing this subject has to a documented incident.



Herman

IEEE Spectrum ran a piece in 2021 by Tekla Perry, titled bluntly, "The iPhone 12 Mini Makes Me Sick, Literally." Soneira was the main interview. It covered hundreds of inquiries the magazine received about visual fatigue, headaches, nausea, and in some cases seizures and photo-triggered epilepsy.



Corn

So this is not a fringe complaint from people who read too many spec sheets.



Herman

And the research gap is its own story. Soneira said manufacturers clam up on this with him, that they don't want to talk about these issues. He also noted there are essentially no published studies on pulse width modulation effects in phones, which he called very odd.



Corn

Odd is a generous word for it.



Herman

It's a case where the people shipping the product know exactly what the behavior is, the affected users are loud, and the literature is nearly empty. That's not an accident.



Corn

There's one more mechanism and I think it's the one that answers Daniel's phone question directly.



Herman

The turn-on threshold. An OLED pixel has a minimum voltage below which it simply doesn't fire. It's off. So there's a hard floor where the emitter either lights at some minimum brightness or goes dark.



Corn

Meaning the lowest brightness setting is not near zero. It's the lowest value the pixel can physically produce.



Herman

LCD backlights have the same problem at minimum voltage. The panel can't ask the light source for less than the light source can give.



Corn

And that's why the minimum brightness on a phone is still too bright. It's not a software oversight. It's a device telling you the truth about its floor.



Herman

Which is the whole episode in one line. Every dimming spec is a conversation about a floor, and the floor is set by physics before it's set by design.



Corn

There's a bridge here that I don't want to skip, because it's the reason any of this matters beyond spec sheet bragging rights.



Herman

The flicker problem is the bridge. It's not just an annoyance. It's the reason the circadian science is relevant to a discussion about driver design.



Corn

So make the case. Because Daniel's claim is that degree of exposure matters more than blue content, and I want to test that rather than nod along.



Herman

The blunt version comes from a lighting product developer who posts on Hacker News. I'm paraphrasing but the phrasing was good, so: you have to reduce both blue light and intensity to avoid suppressing melatonin. Reducing blue light alone might help a little, but it still suppresses melatonin.



Corn

Say the mechanism part again.



Herman

Melatonin levels and circadian phase shifts scale with total irradiance, even when the light is blue depleted. And the punchline he drew from that was, basically, dimming the lights is really effective.



Corn

So a light that shifts color temperature but stays bright is half a tool.



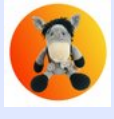
Herman

It's the half of the tool that's easiest to sell, because it's the half you can print on a box.



Corn

What does the peer reviewed side say? Because I want to know whether this is one developer's opinion or something with a spine.



Herman

Joyce, Spitschan and Zeitzer, twenty twenty two, in the Proceedings of the Royal Society B. They found a dose response relationship between light intensity and circadian phase shift. The amount of light drives the circadian response.



Corn

Dose response is the useful phrase. It means more light, bigger effect, along a gradient.



Herman

It does. And I should be honest about the limit of that paper, because it's real. In their flash paradigm they found no dose response for acute melatonin suppression or for alertness. So the picture is more complicated than a clean slogan.



Corn

Meaning intensity governs the phase shifting more clearly than it governs how sleepy you feel an hour later.



Herman

That's how I read it. Same paper, different endpoints, different sensitivity.



Corn

And the counterpoint, because I think we should give it.



Herman

Amdisen and colleagues, twenty twenty two, Chronobiology International. Three hundred and seventeen people, field study, real world daily life. Neither white nor blue light intensity in the morning, evening or night was clearly associated with sleep quality.



Corn

So at ordinary exposure levels, the effect is hard to detect.



Herman

Which is not the same as saying it isn't there. It's saying dose and threshold matter. Somebody sitting in a normally lit room at nine at night is not running an experiment. Somebody staring at a bright screen at midnight might be.



Corn

So the honest framing is that this is a lever with a gradient, not a switch. And the deeper you can pull the lever, the more control you have over where on the gradient you sit.



Herman

Which is exactly the case for wanting granular dimming and not just Kelvin adjustability. If circadian impact scales with irradiance, the amount of light is the variable you'd most want to control precisely. And that's the variable the spec sheet describes worst.



Corn

So Daniel's recommendation gets a scientific spine, and it's not the one the marketing copy would give it.



Herman

It's a better one. You're not buying deeper dimming because it's premium. You're buying it because the thing that affects your sleep is the thing you're now able to set more finely.



Corn

Back to the phone minimum brightness thing, because I want to push on the ambient sensor angle.



Herman

It's the culprit a lot of people miss. The light sensor on a phone or laptop is mediocre at judging what the room actually needs. Manufacturers know this, so they bias toward over brightness. Too bright is annoying. Too dim on a device people need to read is a support ticket.



Corn

So the default drifts upward on purpose.



Herman

And most people never touch it. One comment I saw put it well, that most people still have their screens set at the default brightness level, which is way too bright for anything but direct sunlight.



Corn

The cohort Daniel belongs to is not a cohort of obsessives who hacked a setting. It's a cohort of people who tried to turn it down, found the bottom, and discovered the bottom wasn't low enough.



Herman

That's a hardware floor, not a preference problem. That's the turn-on voltage again.



Corn

Which means the demand for deeper dimming floors in phones and the demand for deeper dimming floors in workbench lighting are the same demand.



Herman

Same population, even. The people who find a screen floor intolerable are the people who go looking for a fixture that can hold a low glow without stuttering.



Corn

Alright. So what's actually changing on the manufacturing side? Because Daniel asked that and I don't want to leave it as a complaint.



Herman

The most interesting answer is hybrid dimming. Soneira's line on this was that some OLED displays already use a combination of analog brightness together with pulse width modulation, which helps.



Corn

Explain why that works, because it's clever.



Herman

Analog handles the low end, where the duty cycle would otherwise be so short that flicker becomes perceptible. Pulse width modulation handles the high end, where it's efficient and the flicker is buried. You use each technique where it's strong.



Corn

It's not one technique beating the other. It's admitting both have a bad zone.



Herman

That's the honest engineering answer. Higher pulse frequencies help too, because they push the flicker above what anyone can sense. Direct current constant current drivers in premium fixtures get you out of duty cycle territory entirely.



Corn

Then the controller side, which I think is the one buyers can actually act on.



Herman

Better driver chips with more bit depth. Ten bit and up instead of eight. That's what turns "stepless" into something real at the floor, because now the bottom one percent gets dozens of steps instead of two.



Corn

Double the bit depth and you quadruple the resolution. Ten bits is a thousand and twenty four steps instead of two hundred and fifty six.



Herman

The smart integrations matter here, Zigbee, Matter, DALI, because they expose finer control curves to whatever's driving the light. But here's the thing Daniel's already onto.



Corn

The LED isn't the bottleneck.



Herman

The controller is. You can have a magnificent emitter and a driver chip that quantizes your beautiful fade into eight visible jumps. The weakest link sets the experience, and the weakest link is usually the cheapest chip in the chain.



Corn

Which is why two lights with identical emitters and identical dimming range specs can feel completely different.



Herman

Why review sites testing for flicker and pulse width modulation have become useful. RTINGS on monitors and TVs, NotebookCheck on laptops, tablets and phones. They're measuring the thing the spec sheet won't state.



Corn

That's the practical takeaway for a listener, and it's not a product recommendation. It's that the specification regime is failing you and you have to get the information from somewhere else.



Herman

Hacker News: Hilbert, what's the actual thing you're sitting on?



Hilbert

What does a hundred to one feel like at the bottom, in a room, with furniture in it?



Corn

That's the question, isn't it.



Hilbert

I've got a cousin. Did lighting for theater and live shows for years. Rigging, focusing, sitting at a desk running a board while somebody on stage argues about their face.



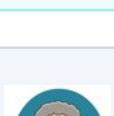
Herman

He's the practitioner.



Hilbert

He told me the real test of a dimmer isn't the range on paper. It's whether the light can hold a candlelit level glow without stuttering or shifting color while it does it. Most cheap dimmers fail that test, and you can see it immediately on stage, because a stage is a place where everyone is looking directly at the light.



Corn

He'd reject fixtures on that basis.



Hilbert

He'd reject fixtures that looked great on paper. He'd pull one out of a rig because the fade wasn't clean. The number on the sheet was fine. The fade was not.



Herman

The curve mattered more to him than the floor.



Hilbert

He said the floor is a claim and the curve is the behavior. A light that only goes to five percent but fades smoothly is a better instrument than one that reaches point one percent in ugly steps. Audiences don't read spec sheets. They watch a fade.



Corn

That's a useful thing to have in your head, because it's the opposite of how the industry sells.



Hilbert

He also kept a lamp in the shop that he'd call the honest one. Old thing, no smart anything. Went down to about four percent and made a clean descent the whole way. He used it to check whether a new fixture was lying to him.



Herman

A reference standard.



Hilbert

A lamp. He said you don't need a meter to know a fade is wrong. You need to watch it next to something that's right.



Corn

What does Daniel's purchasing question actually reduce to, given all that?



Herman

The spec sheet tells you the floor. It doesn't tell you the curve, and the curve is what you perceive. So the industry is measuring the parameter that's easiest to state and not the one that determines whether the thing is pleasant to use.



Corn

The science is pushing from the other side. If intensity is a circadian lever, the buyer wants fine control over intensity, and the market's language for fine control is a word, "stepless," that doesn't mean what it sounds like.



Herman

There's a real gap there. We have a physiological case for granular dimming, a set of well understood engineering obstacles to it, and a measurement regime that describes it badly. That's three things out of alignment and only one of them is likely to move first.



Corn

The research will keep accumulating, because circadian work is going in the direction of dose and intensity rather than blue light alone.



Herman

Pressure will build on manufacturers, though whether they respond by building better drivers or by putting a nicer word on the box is unclear to me. The cheaper path is always the second one.



Corn

Which is why the community testing fills the gap. RTINGS and NotebookCheck and forums full of people who bought the thing and watched the fade.



Herman

That's a strange state of affairs for a mature industry, that the reliable information lives with hobbyists rather than datasheets.



Corn

It happens more often than people admit. And it usually means the datasheet is optimized for a buyer who isn't the person using the light.



Herman

Yeah, I think that's the honest read. The spec exists to sell the fixture. It doesn't exist to describe your evening.



Corn

The open question, then. If the curve matters more than the floor, and no spec sheet tells you the curve, how is somebody supposed to evaluate a light before they own it?



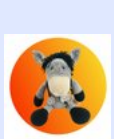
Herman

Watch a fade. Find something you trust, put the two side by side, and trust your eyes over the number on the box.



Corn

That's the whole thing, isn't it. We've got the science pointing one way, the engineering constraints well understood, and a measurement system that doesn't quite capture the part that matters. Somebody's going to have to fix that.



Herman

Thanks to Hilbert Flumingtop for producing the show, and holding the studio together while we argue about lux.



Corn

This has been My Weird Prompts. If you want to send us something, email us at show at my weird prompts dot com. We'll be back soon.



Herman

See you tomorrow.