Episode Book

EPISODE #5681

Why Your Volume Knob Can't Find the Right Spot

PUBLISHED	2 October 2026 at 14:43 UTC
DURATION	21:49
LISTEN	myweirdprompts.com/episode/volume-vs-loudness-explained/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.2
BOOK GENERATED	19 October 2026 at 17:38 UTC

SYNOPSIS

You press the volume button once and it's too quiet. Press it again and it's too loud. Daniel wrote in about exactly this: Android's coarse volume increments, a one-percent-increment app that still didn't help, and a DAC with a single dial that just worked. This episode separates level, volume, and loudness — three terms everyone uses interchangeably and only one of which engineers actually measure. Along the way: why digital attenuation collapses your signal-to-noise ratio at low levels, what dither hides but doesn't undo, why the 1995 Jeff Rowland comparison of digital vs. analog volume control still describes a 2026 phone problem, and how EBU R 128 tried to end the loudness war by fiat. Plus the psychoacoustics of why quiet audio sounds "grating" — and why a hundred volume steps just gives you a hundred wrong answers.

TRANSCRIPT



Corn

There's a specific kind of madness in trying to find the volume that's quiet enough not to wake a sleeping baby and loud enough that you can actually hear the podcast.



Herman

And you can feel it happening in real time. You press the button once, it's too quiet. You press it again, now it's too loud. There's no in-between, and you're sitting there in the dark, jabbing at a rocker like it owes you money.



Corn

Daniel wrote in about exactly this. He's been wrestling with his Android volume control, and the increments are too coarse. He finds an app that does one percent increments, and somehow it still doesn't fix it. He gets it close, then it's just slightly too loud and too grating. Then he goes to a proper speaker with a DAC, one dial, zero to a hundred, and it just works. You dial in the exact right level and go back to listening.



Herman

Same signal, same ears, totally different experience.



Corn

And that's what he wants to dig into. What actually is volume, versus loudness? Because anyone who's worked with audio editing knows those words aren't synonyms at the technical level. He's asking about normalization, EBU broadcast loudness standards, and the whole chain of mastering, rendering and amplification that sits between the file and your ear.



Herman

And why that chain produces moments where you simply cannot get the right sound, no matter how many steps you give yourself.



Corn

There's a lot in there. So let's start by separating the words, because I think the whole problem lives in that gap.



Herman

Three terms that everyone uses interchangeably, and only one of them is what engineers actually measure. Level is the measurable quantity. It's sound pressure in pascals, it's volts if you're looking at an electrical signal, and in digital audio it's dBFS, decibels relative to full scale. Zero dBFS is the loudest thing a digital system can carry. Everything below it is negative.



Corn

So level is the physical fact.



Herman

Level is the physical fact. Volume is a legacy word. It comes from the knob on a radio set, and it's not standard vocabulary in audio engineering at all. When somebody says volume, they usually mean the quantity of sound in the room, which is really a statement about the listener's experience, not about the signal.



Corn

And loudness is the perception itself.



Herman

Loudness is subjective, full stop. The ANSI definition is that it's the attribute of auditory sensation by which sounds can be ordered on a scale from quiet to loud. It depends on intensity, on frequency, on duration. Two signals can sit at exactly the same peak level and have completely different loudness.



Corn

Which is the entire loudness war in one sentence.



Herman

That's the mechanism. You compress the dynamics, you bring the quiet parts up toward the loud parts, then you apply make-up gain so the whole thing rises. The peaks stay where they were, but the average level climbs. You've added loudness without adding level. That's how a record from two thousand eight can feel exhausting next to one from nineteen eighty.



Corn

So the thing Daniel is chasing with his finger on a volume rocker is loudness, and the thing the phone is actually adjusting is level.



Herman

And the phone is adjusting it badly, in a way that has almost nothing to do with how many steps it gives you. Which is the part I find interesting.



Corn

Start with the steps, though, because Daniel's specific complaint is about the coarseness.



Herman

Android's default media volume is fifteen steps. In-call volume is seven. Those numbers aren't a design decision made in a meeting about user experience. They're compiled into the system image as build properties, read at boot by the audio service. If you're rooted, you can edit them. Thirty steps is the common tweak. Samsung ships a SoundAssistant app that gives you fine control without root. And there's a developer options setting to disable absolute volume that changes how Bluetooth headphones negotiate levels.



Corn

So the number is real, and it's changeable.



Herman

It's changeable, and here's the thing. Raising it from fifteen to thirty does not fix Daniel's problem. Neither does an app that gives him a hundred. And I think that's the heart of the episode, because the intuition is that more steps means finer control means better outcome. That intuition is wrong.



Corn

Why is it wrong?



Herman

Because the damage happens when you turn the thing down. Digital attenuation scales the signal and leaves the noise floor exactly where it was. So as you reduce the level, your signal-to-noise ratio collapses. Analog controls attenuate the signal and the noise together, so the ratio holds steady all the way down.



Corn

Give me the numbers on that, because that's the kind of claim people argue about.



Herman

Sixteen-bit audio, forty-four point one, a DAC with a noise floor around minus one hundred and four decibels. At minus ten dBFS, analog gives you ninety-six decibels of signal-to-noise. Digital gives you ninety-four. Basically a wash, nobody could tell. By minus forty dBFS, analog is still at ninety-six, and digital has fallen to sixty-four. That's a thirty-two decibel gap, and thirty-two decibels is not subtle.



Corn

So the quiet listening scenario, the one where you're trying not to wake a baby, is exactly the scenario where the digital path is at its worst.



Herman

The use case that made Daniel reach for the volume control is the use case that digital attenuation handles most badly. Every six decibels of attenuation costs you roughly one bit of resolution. Sixteen-bit gives you ninety-six decibels of dynamic range to start with, twenty-four gives you a hundred and forty-four. Turn a sixteen-bit signal down forty decibels and you have given away more than six bits of your resolution to arithmetic.



Corn

Where does the resolution actually go?



Herman

It gets truncated. You're multiplying every sample by a factor less than one, and the result is almost never an exact integer at sixteen or twenty-four bits. The remainder gets rounded off, and the rounding pattern is not random, it's correlated with the signal, which is what makes it audible as distortion rather than hiss. Dither is the fix. You add a tiny amount of low-level noise so the truncation error decorrelates and turns into a gentle hiss instead of a gritty edge. It works. But dithering hides the loss, it doesn't undo it. The resolution is still gone.



Corn

So the phone is doing arithmetic on a sixteen-bit stream, throwing away bits as it goes, and then handing the result to an amplifier that's just faithfully reproducing a degraded signal.



Herman

And the DAC dial is doing something structurally different. Analog attenuation brings the noise down with the signal, so the ratio stays intact.



Corn

Which raises the question of why anyone ever put the volume control in the digital domain in the first place.



Herman

Cost, mostly. A digital volume control is free. It's a multiply. An analog volume control is a physical component, and a good one is expensive. There's a nineteen ninety-five paper from Jeff Rowland Design Group comparing DAC-based digital volume against the CS3310 analog volume chip, and the conclusion was that linear-step digital control is very coarse and inaccurate at high degrees of attenuation. That's a thirty-year-old finding and it describes Daniel's problem exactly.



Corn

There's a rebuttal to that, though, and it's a good one.



Herman

There is. Thirty-two-bit DACs with internal digital volume control change the arithmetic completely. An ESS Sabre chip doing the attenuation inside its own data path has enough headroom that you can turn it down substantially without ever hearing quantization error. The catch is that it only works if the volume control has access to the DAC's internal pipeline. If the attenuation is happening in the media player, three layers up the stack, then you bought a good chip and never used it.



Corn

So the answer to analog versus digital is: it depends where the multiply happens.



Herman

That's the honest version. Digital volume done at the DAC is excellent. Digital volume done before the DAC is where the thirty-two decibel gap lives.



Corn

Okay, so that's the engineering. Here's what I keep circling back to, though. Even when you fix the gain structure, even with infinite steps, the sleeping baby problem doesn't go away. Because Daniel said something specific. He said he gets it close and then it's grating.



Herman

Grating is the interesting word.



Corn

Grating is a perceptual description, not a level description. Something can be quiet and still be grating.



Herman

That's the psychoacoustics, and this is the part where I have to be honest about the limits of what I know. There's a well-established phenomenon here, the equal-loudness contours, the Fletcher-Munson work. Human hearing is not flat. At low listening levels, bass and treble fall away faster than the midrange. So as you turn something down, the tonal balance shifts. You don't just get a quieter version of the same sound, you get a thinner one.



Corn

Which means there's a whole region of the volume range where the sound is technically present but perceptually hollow.



Herman

And thin, midrange-heavy sound is exactly what reads as grating. The ear is most sensitive in the two to four kilohertz region. That's the band that carries intelligibility and also the band that carries harshness. So a podcast turned down low, with the bass gone, leaves the midrange exposed. You hear the consonants and the sibilance and not much else.



Corn

That's a good explanation for the thing Daniel couldn't articulate. He couldn't say why it felt wrong, just that it did.



Herman

I want to flag that I'm connecting two things there rather than citing a study that measured it on Android specifically. The contour effect is solid. Whether Android's specific step curve interacts with it in a particular way is something I don't have a source for. There's a real gap in what's documented. The step count is confirmed, the curve shape isn't.



Corn

Fair. But it explains something else, too. Why the app with one percent increments didn't save him.



Herman

It couldn't. He was hunting for a point on a curve that was never going to feel right at that level, because the curve itself is the problem. A hundred steps gives you a hundred wrong answers instead of fifteen.



Corn

That's grim and I believe it completely.



Herman

There's a second thing going on, too, and it's about the interface rather than the signal. But I want to get to the standards first, because the standards are the part of this story where somebody actually tried to solve it properly.



Corn

The loudness war, and the attempt to end it by fiat.



Herman

EBU R 128. First published in August twenty ten, put together by the PLOUD group, led by Florian Camerer at ORF. What it did was change what broadcasters normalize to. Before that, everything was peak normalization. You set your ceiling at a certain level and let the average fall wherever it fell, which meant a quiet drama and a loud commercial could peak at the same place and be twenty decibels apart in perceived loudness.



Corn

So the ad breaks were the loudness war, in broadcast form.



Herman

The ads were the whole reason it got fixed. Nobody in television was going to keep fielding complaints about how the commercials were deafening. R 128 moved the target from peaks to loudness. The number is minus twenty-three LUFS, integrated across the whole programme, with a tolerance of half a LU, one LU for live, and a maximum true peak of minus one dBTP.



Corn

Break those units down, because they get thrown around a lot.



Herman

LUFS is Loudness Units relative to Full Scale. In the ITU standard it's called LKFS, same thing, different letters. A LU is a relative unit, and one LU equals one decibel. LRA is loudness range, which tells you how much the loudness moves across a programme. dBTP is true peak, which accounts for inter-sample peaks that a sample-accurate meter will miss.



Corn

And the measurement isn't just an average.



Herman

It's a specific algorithm. You apply K-weighting, which is a frequency curve approximating how human hearing responds, applied to every channel except the low-frequency effects channel, because you don't want the subwoofer driving the measurement. Then you meter in three windows. Momentary, four hundred milliseconds. Short-term, three seconds. And integrated, across the whole programme.



Corn

And the integrated figure has gates.



Herman

Two. An absolute gate at minus seventy LUFS, so silence and near-silence don't drag the average down. And a relative gate ten LU below the current integrated value, so a quiet passage doesn't get to redefine the whole measurement. It's a elegant piece of engineering, and it's why a commercial break doesn't blow your head off on a properly compliant broadcast.



Corn

That's a real fix. So why doesn't it help Daniel with his phone?



Herman

Because R 128 governs the production and distribution end. It makes sure the content arrives at a predictable loudness. It says nothing about the gain structure of the device that plays it back. Daniel's problem is downstream of every standard in this conversation.



Corn

And streaming has its own layer of this now.



Herman

Every platform normalizes to its own target. Spotify lets you pick between minus eleven, minus fourteen and minus nineteen LUFS. YouTube, Tidal and Amazon Music normalize to minus fourteen. Apple Music sits at minus sixteen with Sound Check. Qobuz uses minus eighteen. So the same master plays back at a different loudness depending on which app you opened.



Corn

Which is why a podcast sounds different on two services.



Herman

And why the loudness war is, depending on who you ask, over or not over. Wikipedia's phrasing is that streaming normalization arguably put an end to it. But there was a two thousand twenty-three paper on music de-limiter networks that treats it as an ongoing phenomenon, still causing ear fatigue and hearing loss. Both can be true, I think. Peak-limited loudness stopped paying off the way it used to, but the incentive to sound louder than the next track in a playlist did not disappear.



Corn

The gate moved. It didn't close.



Herman

And there's fresh work in this space. A paper from last November applies equal-loudness contours as a loss function for speech enhancement. So the psychoacoustics Daniel ran into by accident at two in the morning is an active research area for people building speech models.



Corn

Give me the practical version. If Daniel is producing a podcast, or if he's just trying to not wake the baby.



Herman

If you're producing, master to minus fourteen LUFS for streaming, minus twenty-three for broadcast, and watch your true peak as carefully as your integrated loudness. The peak is what clips when the platform's own processing nudges your track up.



Corn

If you're listening.



Herman

If you're listening, the advice is unsatisfying. Don't hunt for the exact step. The step count was never your problem. Keep the source material at a healthy level going into the device, so you're attenuating as little as possible, and if you care about quiet listening, the analog control downstream is doing more work than any app on the phone.



Corn

Which is a slightly annoying conclusion for a podcast about technology.



Herman

It's an annoying conclusion for me too. I'd prefer the software to win.



Corn

There's one more piece, though, that we haven't touched, and it's the one Daniel actually led with. Why the dial feels different in the hand.



Herman

The tactile feedback thing. I looked for a technical source on that and didn't find one. The documented advantage of the analog control is the SNR preservation and the bit depth, not the feel of the knob. But I think there's something real there that just doesn't show up in the measurements, and I'd rather say that plainly than pretend I have a paper.



Corn

Then it's a good thing we have someone at the desk who's actually owned one of these.



Hilbert

Yamaha CD-S2100. Twelve hundred dollars used, and a hundred and thirteen and a half decibels of attenuation in half-decibel steps, motorized knob, and I sold it in the spring of eighty-six because the furnace needed replacing.



Corn

Half-decibel steps.



Hilbert

The furnace was eleven hundred. The man told me the old one had maybe two winters left in it. I told him two winters was two more than the stereo was going to get me, and I wrote the check.



Herman

So the DAC went.



Hilbert

The DAC went. And the thing about that knob was that it clicked, so you knew where you were. Every half a decibel was a click you could feel through your fingers, and the position told you your gain. You did not have to press anything and hope.



Corn

You could find the same level twice.



Hilbert

You could find it without looking. My neighbor's kid was sleeping on the other side of that wall, and I used to do the mix at two in the morning, and I never woke him once. You know what I was doing at two in the morning? I was trying to hear the low end on a cassette transfer without the whole house knowing about it.



Herman

The phone gives you none of that.



Hilbert

The phone gives you a button. You press the button, something changes, you have no idea by how much, so you press it again, and now it's worse. That is not a control. That is a slot machine.



Herman

Confidence as a piece of the signal chain. That's the part I hadn't thought about.



Hilbert

It's not the signal chain. It's you. If you don't believe you can get back to where you were, you can't settle, so you keep fiddling, and every time you fiddle you're further from where you started. The half-decibel detents meant I could settle. That's worth more than the thirty-two decibels of signal-to-noise, if I'm honest, because a man who's settled stops touching the knob.



Corn

The motorized knob. Was that the thing that made it half-decibel?



Hilbert

No. The motor was so the remote could move it. The steps were the chip. The chip was a Cirrus, and the whole point of it was that it was analog, so the noise came down with the signal. Same reason the digital ones are bad, opposite direction.



Herman

Which is exactly the CS3310 family we were talking about. You had the good version of the thing.



Hilbert

I had the good version of the thing, and I sold it to pay for a furnace, and I have priced that knob approximately once a year for forty years. I put the levels on your bumper at minus eighteen, by the way. Herman was pushing into it on the last record.



Herman

I was not.



Hilbert

You were. Two hundred milliseconds of it. I'll clean it in the edit.



Corn

Okay. So the episode Daniel asked for has a slightly different shape than I expected when we started.



Herman

It does. Because the engineering answer and the perception answer converge on the same place, and it's not the place I would have guessed. The thing Daniel is missing isn't precision. It's confidence, and it's a tonal balance that survives being turned down. The phone can't give him the second one because the physics won't allow it, and it can't give him the first one because a rocker with no detents is a control that refuses to tell you where you are.



Corn

The DAC dial gives him both.



Herman

It gives them to him without any of the things we normally credit. No better DAC, no better amplifier, just an attenuator that steps in analog, in half-decibels, in a position you can feel.



Corn

Here's a thing I didn't expect to be thinking about at the end of this. There's a paper from last year applying equal-loudness contours to speech enhancement. The research is heading toward devices that model your hearing and adapt. You could imagine a phone that knows what you're listening to, knows the room, knows your ears, and sets the loudness for you.



Herman

You could. And I'd want to know how it handles the case where you disagree with it. Because the entire sleeping baby problem is a negotiation between what you want to hear and what the room will permit, and I'm not sure I trust a model to know which side of that line I'm on.



Corn

That's the interesting question, though. Whether the fix is more steps or fewer decisions.



Herman

Given how this episode went, I'd bet on the detents.



Corn

Before we go, one thing from the research that didn't make it in. The Wikipedia entry on the loudness war has a specific line about streaming normalization, and it says the practice arguably put an end to it. It's the word arguably doing all the work in that



Corn (continued)

sentence. A decade of engineers fighting over how loud a master should be, and the resolution is a qualifier in a wiki article.



Herman

The two thousand twenty-three de-limiter paper treating it as ongoing. Both of them are right, which is the most engineer thing that could possibly have happened.



Corn

That's the show. Thank you to our producer, Hilbert Flumingtop. I'll be thinking about the click of that knob for a while.



Herman

Same. A nine-hundred-dollar detent is going to live in my head rent-free.



Corn

This has been My Weird Prompts. If you want to send us something to argue about, email us at show at my weird prompts dot com.



Herman

We'll be back soon.



Corn

Take care of your ears.