MY WEIRD PROMPTS

Podcast Transcript

EPISODE #212

The Authenticity Crisis: Proving You’re Real in 2046

Published January 10, 2026 • Runtime: 23:15

https://myweirdprompts.com/episode/ai-authenticity-crisis-future/

EPISODE SYNOPSIS

As AI-generated content becomes indistinguishable from reality, we are entering a fundamental crisis of trust where "seeing is believing" no longer applies. In this episode of My Weird Prompts, Herman and Corn dive deep into the technical and philosophical battle for truth over the next twenty years. They explore the rise of "controlled capture" hardware, the cryptographic signatures of the C2PA, and the controversial emergence of biometric "Proof of Personhood" systems like Worldcoin. The discussion moves beyond simple deepfakes to examine the terrifying possibility of "Reality as a Service," a future where digital authenticity is a paid luxury and the "Dead Internet Theory" becomes a daily reality for the unverified. From the "Authenticity Renaissance" of raw, imperfect media to the concept of "Social Mining" in physical spaces, Herman and Corn map out the high-stakes arms race between synthetic perfection and human imperfection. Join us for a look at how we will safeguard our identities in an era where the mouse has a jetpack and the truth has a subscription fee.

DANIEL'S PROMPT

Daniel

With the rapid advancement of AI and its ability to create deepfakes or doctor almost any digital content, what will the landscape look like in 10 to 20 years for individuals who need to prove that a digital artifact—like a photo, video, or chat—is real, or even prove their own humanity?

TRANSCRIPT

Corn

Hey everyone, welcome back to My Weird Prompts. I am Corn, and I am sitting here in a very rainy Jerusalem with my brother.

Herman

Herman Poppleberry, at your service. And yeah, the weather is definitely matching the mood of the prompt our housemate Daniel sent over this morning. It is a bit of a heavy one, but honestly, it is something we have been skirting around for a while now.

Corn

It really is. Daniel was asking about the landscape of truth, basically. With AI getting so good at faking everything, how do we prove anything is real in ten or twenty years? Not just proving a photo is real, but proving that you, the person sending the photo, are actually a biological human being and not just a very convincing script.

Herman

It is the authenticity crisis. We have talked about the technical side of this in bits and pieces, like back in some of our recent episodes when we were looking at deep packet inspection and how networks try to verify traffic. But Daniel’s question goes much deeper into the personal and social level. It is about the fundamental erosion of "seeing is believing."

Corn

Exactly. I mean, we are already seeing the early signs of this in 2026. You cannot trust a video call anymore without some kind of secondary verification. But looking ten years out, to twenty thirty-six, or twenty years to twenty forty-six... that is where it gets really weird.

Herman

It does. And I think to understand where we are going, we have to look at the tools being built right now. Daniel mentioned an app he found that adds metadata and calculates hashes to prove a photo is real. He is likely talking about something like Truepic or the Oaro Media platform. These are essentially trying to create what we call "controlled capture."

Corn

Right, and for those who are not familiar, the idea there is that the verification happens the millisecond the light hits the camera sensor. It is not about taking a photo and then trying to prove it is real later. It is about the hardware itself saying, "I am a certified sensor, and I am signing this data right now."

Herman

Precisely. We are starting to see this move from labs and pilots into more mainstream products. The Coalition for Content Provenance and Authenticity, or C2PA, has been iterating on its specification, and there is active work and discussion around how to extend content credentials toward live or near-real-time media. It is not a fully solved problem for live streaming yet, but the direction of travel is clear: the goal is that if you are watching a broadcast, you will eventually be able to see a little content-credential icon that proves the video feed has not been tampered with between the camera and your screen.

Corn

But Herman, here is the thing that bothers me. If we move to a world where everything has to be "signed" by the hardware to be considered real, what happens to the average person who does not have a high-end, provenance-aware phone? Do they just become "fake" by default?

Herman

That is the big risk. We are looking at a potential "authenticity divide." If you look at newer flagship phones like Google’s latest Pixel models and Apple’s newest iPhones, you can already see a trend toward secure enclaves and hardware security modules that could be used for things like C2PA-style signing. When you take a photo, future versions of this tech could embed a cryptographic manifest that lists the time, the location, and whether generative AI was used to alter the pixels. But if you are using an older device, or a budget device that does not have that kind of secure hardware, your photos might be flagged as "unverified" by social media platforms or even in legal disputes.

Corn

It feels like we are moving toward a tiered reality. You have the "verified" layer where the wealthy and the tech-savvy live, and then the "synthetic ocean" where everyone else is. Daniel’s prompt mentioned proving your own humanity, too. That is the part that really feels like a science fiction dystopia. How do I prove I am me in twenty years?

Herman

Well, look at what Worldcoin is doing right now. They recently put out a roadmap update on World ID and its next major iteration. They want you to scan your iris with one of their "Orbs" to prove you are a unique human. They use techniques like secure hardware and zero-knowledge proofs to try and keep that data private, but the core idea is that your biology becomes your digital passport.

Corn

I have always found the Orb a bit creepy, honestly. But I can see the logic. If an AI can pass the Turing test, if it can write code better than I can, and if it can generate a video of me saying things I never said, then my biological signature might be the only thing left that is "unforgeable." Or at least, very hard to forge.

Herman

It is the "Proof of Personhood" problem. In ten years, I suspect we won't just be signing into websites with a password or even a passkey. We will be using zero-knowledge proofs. These are cryptographic methods where you can prove you are a verified human without actually revealing who you are. So, the website gets a "yes" or "no" from a trusted identity provider, but they don't get your name or your biometric data. It is like showing a bouncer a card that says "I am over twenty-one" without showing them your name or address.

Corn

But that assumes we trust the identity provider. We are basically just moving the needle of trust from "I trust my eyes" to "I trust this massive corporation or decentralized protocol." It doesn't actually solve the problem of truth; it just centralizes the verification of it.

Herman

That is a great point, Corn. And it leads to a second-order effect that most people are not thinking about. If truth becomes a service provided by companies, then truth has a subscription fee. If you want your emails to be marked as "human-written" so they don't get filtered into the spam folder of the future, you might have to pay for a verification service. We are talking about the "Dead Internet Theory" becoming a reality for anyone who isn't paying for a seat at the table.

Corn

That is a terrifying thought. "Reality as a Service." If you don't pay your twenty dollars a month to the Authenticity Bureau, nobody believes a word you say online. But let's flip it for a second. What if the fakes get so good that the "verified" signals actually become the target? If I am a sophisticated attacker in twenty thirty-five, I am not trying to make a better deepfake; I am trying to hack the signing key for your camera or phone.

Herman

Oh, absolutely. The arms race never ends. Once you create a lock, someone creates a lockpick. If the content-credential icon becomes the gold standard for truth, then stealing a private key from a camera manufacturer becomes the ultimate heist. We have seen in other industries that when signing keys leak—like past leaks of driver certificates or console firmware keys—it opens the door to malware or unauthorized firmware that still looks "official." The same risk would apply if a camera maker’s keys were ever compromised. You could generate a completely fake war or a fake political scandal and give it the "Verified" badge.

Corn

It reminds me of the "Portable Fortress" discussion we had last week in one of our recent episodes. We talked about moving your network and your security with you. In the future, you might need to carry your own hardware-based identity vault everywhere you go, just to interact with the world.

Herman

Exactly. And I think we need to address a common misconception here. A lot of people think that "deepfake detection" is the answer. You know, the AI that looks for the weird blinking patterns or the distorted earlobes. But the truth is, that is a losing battle. By twenty thirty, detection software will likely be outmatched because the generative models will be trained specifically to bypass the detectors. It is a cat-and-mouse game where the mouse eventually gets a jetpack.

Corn

Right. It is the "Generative Adversarial Network" problem. The faker and the detective just keep making each other better until the faker is perfect. So, the only solution isn't detection; it is provenance. It is knowing where the file came from, not what is in it.

Herman

Yes! Provenance is the key. It is the digital equivalent of a chain of custody for evidence. But here is where it gets really interesting, and maybe a bit more hopeful. There is a growing trend right now—some people are calling it the "Authenticity Renaissance" or the "Rawness Movement."

Corn

I have seen this. It is that idea that because AI makes everything look "perfect" and "polished," humans are starting to crave the opposite. Like that "No-Filter" movement on social media, but on steroids.

Herman

Exactly. There have been tech columns recently arguing that "AI makes polish cheap." If you can generate a perfect, cinematic video for almost nothing, then a perfect, cinematic video has no value. But a shaky, slightly out-of-focus, raw video with background noise? That feels human. It is the "uncanny valley" in reverse—we are looking for the flaws to find the soul.

Corn

So, the proof of humanity isn't a cryptographic hash; it is the fact that I messed up the lighting?

Herman

In a way, yes! We are seeing creators deliberately lean into "imperfection" as a signal of authenticity. It is like the digital version of a vinyl record. People want the "hiss" and the "crackle" because it proves it was made in the physical world.

Corn

I love that. But I worry that AI will just learn to fake the crackle. It will learn to add the "shaky cam" and the "mumbled words" to make itself look more human. In fact, we already have models that do that.

Herman

They do. But there is a level of "contextual rawness" that is very hard to fake. For example, if I am recording a video in our kitchen in Jerusalem, and you can see the specific way the light hits that one cracked tile we have been meaning to fix, and then I walk outside and you see the neighbor's cat—the one with the notched ear—those are hyper-local, high-entropy details. An AI can't know those things unless it has been surveilling our house in real-time.

Corn

So, "Proof of Humanity" becomes "Proof of Locality"?

Herman

Exactly. It is about being tethered to a physical space. In ten or twenty years, I think our digital identities will be much more closely tied to our physical presence. We might see "Proof of Proximity" where two people’s phones exchange a cryptographic handshake when they meet in person, which then boosts their "Humanity Score" on social networks. It is like a digital vouching system.

Corn

That is a fascinating thought experiment. Imagine a world where your "Humanity Score" goes up every time you hang out with other verified humans in the real world. It turns social interaction into a form of mining for authenticity.

Herman

"Social Mining." I like that. It sounds a bit like a Black Mirror episode, but it might be the only way to stay ahead of the bots. If the bots can't easily inhabit the physical world—well, until we have humanoid robots everywhere, which is a whole other episode—then the physical world remains our "Safe Harbor."

Corn

We actually touched on the robot side of things in one of our other episodes, "Predictive Motion." Once the bots can walk and move like us, even the physical world becomes suspect. But let's stick to the digital side for now. Daniel’s question was about individuals needing to prove a digital artifact is real. Think about a legal case. If I am in a car accident in twenty thirty-six, and I take a photo of the damage, how do I make sure the insurance company doesn't claim I "AI-generated" the dent?

Herman

That is where the app Daniel mentioned comes in. Truepic Vision, for example, is already being piloted and used by some insurance companies for exactly that. When you use their "Controlled Capture" tech, the app doesn't just take a photo. It records the accelerometer data to prove the phone was moving in a certain way. It records the GPS, the weather data, the cell tower IDs, and it bundles all of that into a tamper-evident package. If you try to change one pixel, the whole cryptographic seal breaks. It is basically a digital notary inside your pocket.

Corn

So, for the listener who is worried about this, the practical takeaway is: start looking for these "Provenance-Aware" tools. Don't just use a generic camera app for important things.

Herman

Exactly. And keep an eye on the "Content Credentials" icon. It often appears as a small stylized "CR"-type badge. More and more websites are starting to show it. If you click it, you can see the "manifest"—it will show you if the image was edited in Photoshop, if it was generated by an AI, or if it came directly from a verified sensor.

Corn

But what about a chat exchange? Daniel mentioned a WhatsApp chat. Those are notoriously easy to fake. You can find "Fake Chat Generators" online in two seconds.

Herman

Chats are much harder because they are text-based. Text has very low "entropy" compared to a photo. But we are seeing the emergence of "Verifiable Logs." Imagine a version of a messaging app where every message is written to a private, encrypted ledger. If you want to prove a conversation happened, you can share a "view key" for that specific thread with a judge or a third party. It uses the same tech behind some of the more advanced blockchains, but for your D-Ms.

Corn

It feels like we are losing the "casualness" of the internet. Everything has to be logged, signed, and notarized just to be believed. Do you think we will miss the "Wild West" era where we could just... exist online without a "Humanity Score"?

Herman

I think we already miss it, Corn. But the price of the Wild West was that it eventually got overrun by bandits—or in this case, bots and bad actors. We are moving into the "Fenced Garden" era of the internet. The fences are the cryptographic protocols that keep the fakes out.

Corn

It is a trade-off. Security versus freedom. Authenticity versus anonymity. If I have to prove I am a human to post a comment on a blog, then I can't be anonymous anymore.

Herman

That is the biggest tension. Anonymity is a vital part of a free society, especially for whistleblowers or people in oppressive regimes. If we mandate "Proof of Humanity" via biometrics, we might accidentally destroy the ability to speak truth to power anonymously.

Corn

Unless... and this is where your nerdiness comes in, Herman... unless we use those Zero-Knowledge Proofs you mentioned.

Herman

Exactly! That is the "Aha!" moment. With Zero-Knowledge Proofs, I can prove "I am a human who lives in Jerusalem and is over eighteen" without saying "My name is Herman Poppleberry and here is my ID number." It is the "Holy Grail" of digital identity. It gives us the "Proof of Humanity" without the "Loss of Privacy."

Corn

Is anyone actually building that right now?

Herman

Yes. The Worldcoin project I mentioned is trying to do it, and there are several other protocols like Sismo and Polygon ID that are working on "Privacy-Preserving Identity." It is still early days, but in ten years, I think that will be the standard.

Corn

Okay, let's look at the twenty-year horizon. It is twenty forty-six. AI is a billion times more powerful than it is today. What does a "real" photo even mean then? If an AI can simulate the physics of light so perfectly that it is indistinguishable from a camera sensor, does the concept of a "photo" even exist?

Herman

That is a deep one. I think we might see a return to physical artifacts for the most important things. Maybe the "Birth Certificate" of the future isn't a digital file; maybe it is a physical, analog object that is hard to replicate. Or maybe we move toward "Biometric Anchoring."

Corn

Meaning?

Herman

Meaning every digital file you create is cryptographically "tethered" to your unique biological signature. So, the file isn't just "real"; it is "yours." If I see a video of you, my device checks if it is anchored to your biological key. If it isn't, my phone just shows a big red warning: "Synthetic Identity Detected."

Corn

It is like a digital aura. You carry your "truth" with you.

Herman

I love that term. A "Digital Aura." It is a beautiful way to think about it. It is not just about the data; it is about the "spirit" of the creator being embedded in the file.

Corn

But man, the energy requirements for all this signing and verifying must be insane. If every photo, every chat, every email has to be cryptographically notarized...

Herman

It is a lot of compute, for sure. But as we discussed in our episode on hardware, our chips are getting more efficient at these specific tasks. We have dedicated crypto and security cores in our devices now that can do these signatures in microseconds with minimal battery drain.

Corn

Still, it feels like a lot of work just to get back to where we were in the nineteen-nineties, where you could just look at a photo and know it was a photo.

Herman

(Laughs) Welcome to the future, Corn! We have to build a billion-dollar infrastructure just to recreate the feeling of common sense.

Corn

It is funny because it is true. But honestly, I think Daniel’s "Doom's Day" prediction is a bit too dark. Humans are incredibly good at adapting. We developed a "BS detector" for the printing press, we developed it for television, and we are developing it for AI.

Herman

I agree. We are in the "Confusion Gap" right now. The technology has outpaced our social norms. But in ten or twenty years, checking the provenance of a file will be as natural as checking the "HTTPS" lock in your browser address bar. We won't even think about it.

Corn

I hope you are right. Because the alternative is a world where we all just stop believing anything, and that is a very lonely place to be.

Herman

It is. But that is why we do this show, right? To dig into the weirdness so it doesn't feel so overwhelming.

Corn

Speaking of the show, we should probably wrap this up before we go too far down the rabbit hole. Daniel, thanks for the prompt. It definitely gave us a lot to chew on.

Herman

Yeah, thanks Daniel. And to everyone listening, if you are finding this useful—or if it is just keeping you company while you hide from the rain—we would really appreciate it if you could leave us a review on Spotify or whatever podcast app you are using. It actually helps a lot more than you might think.

Corn

It really does. It helps the "human" signals stand out in the sea of AI-generated content!

Herman

(Laughs) Exactly! We need to boost our "Humanity Score."

Corn

You can find all our past episodes, including the ones we mentioned today, at myweirdprompts.com. We have got an RSS feed there, and a contact form if you want to send us your own weird prompts.

Herman

And we are on Spotify, obviously. Until next time, stay curious and keep checking those credentials.

Corn

This has been My Weird Prompts. We will catch you in the next one.

Herman

Bye everyone!