

MY WEIRD PROMPTS

Podcast Transcript

EPISODE #87

The \$100 Million Giveaway: Why Big Tech Opens Its AI

Published December 23, 2025 • Runtime: 24:09

<https://myweirdprompts.com/episode/open-weights-vs-proprietary-ai/>

EPISODE SYNOPSIS

In this episode of My Weird Prompts, Herman Poppleberry and Corn the Sloth tackle a baffling question from their housemate Daniel: Why are companies like Meta and Mistral spending hundreds of millions of dollars to build massive AI models, only to release the "blueprints" for free? From the \$100 million training costs of Llama 3 to the strategic maneuvers of Mark Zuckerberg, the duo explores the hidden business logic behind "open weights." Is it a play for developer mindshare, a clever way to recruit top talent, or a defensive move against the closed gardens of OpenAI and Google? Herman and Corn debate the security risks of decentralized AI versus the dangers of "security through obscurity," while also touching on the "no moat" theory that suggests the open-source community might be eating the lunch of the tech giants. Grab a snack and join the conversation as they decode the trillion-dollar chess game of the AI industry.

DANIEL'S PROMPT

Daniel

We have seen a significant amount of activity in the world of open source during the AI era so far. Given the major investments and teams required to bring these models to market, why is there so much open-sourcing in AI? What motivates major AI labs to invest so much money in training powerful models only to give them away for free, and do you expect to see more open-source activity versus proprietary models moving forward?

TRANSCRIPT

Corn

Welcome to My Weird Prompts, the show where we take the strange, the complex, and the downright confusing ideas and try to make some sense of them from our living room here in Jerusalem. I am Corn, your resident sloth and professional ponderer, and I am joined as always by my brother, Herman Poppleberry.

Herman

That is Herman Poppleberry, the donkey with the data, at your service. It is good to be here, Corn. I have been diving deep into the research for this one because our housemate Daniel sent us a prompt that is actually pretty baffling when you first look at it. He was asking why on earth these massive companies are spending hundreds of millions of dollars to build artificial intelligence models just to give the blueprints away for free.

Corn

Yeah, it does seem a bit counterintuitive, right? If I spend six months building a really nice treehouse, I am probably not going to just hand the keys to every person who walks by on the street. But in the world of AI, we are seeing Meta, Mistral, and even Google to some extent, just putting their hard work out there. Why would they do that?

Herman

Well, that is the big question. We are talking about open source versus proprietary models. You have got the closed-off gardens like OpenAI and Anthropic on one side, and then you have this massive movement of open weights models on the other. It is a strategic chess game, Corn. It is not just about being nice or being a digital hippie. There are billions of dollars in play here.

Corn

I mean, I want to believe they are just being nice. Maybe they just want the world to have cool tools?

Herman

Mmm, I really don't think so. I mean, look at Meta. Mark Zuckerberg isn't exactly known for doing things just out of the goodness of his heart without a business angle. When they released Llama three back in April of two thousand twenty-four, it was a massive shift. They spent an estimated hundred million dollars or more just on the compute power to train that thing. To just give that away? There has to be a colder, more calculated reason.

Corn

Well, before we get into the cold, calculated stuff, let's look at what open source actually means in this context. Because I heard some people saying Llama isn't actually open source in the traditional sense.

Herman

You are actually right to question that. See, traditional open source, like Linux or the Python programming language, means you have the source code and you can do whatever you want with it. With these AI models, many people prefer the term open weights. You get the finished brain of the AI, but you don't necessarily get the data it was trained on or the exact recipe they used to cook it. But for the average developer, it is close enough because they can run it on their own hardware.

Corn

Okay, so if I am a developer, I love this because it is free. But back to the why. If I am Meta, and I give this away, am I not just helping my competitors? Like, if I give a smaller company Llama three, they can build a product that competes with Facebook, right?

Herman

That is exactly what people thought at first. But look at it from the perspective of standard setting. If everyone uses Meta's architecture, then Meta defines how the industry works. They become the gravity at the center of the ecosystem. If all the developers are building tools that work best with Llama, then Meta has a huge influence over the future of the technology. It is a play for developer mindshare.

Corn

I don't know, Herman. That feels like a very expensive way to get people to like you. A hundred million dollars for mindshare? I could buy a lot of friends for a lot less than that. There has to be something more tangible.

Herman

There is. Think about the talent. The best AI researchers in the world don't want to work in a black box where they can never show their work to their peers. By being open, Meta attracts the top tier scientists who want to publish papers and see their work used by millions. Plus, you have thousands of developers outside the company finding bugs and optimizing the code for free. It is like having a global department of unpaid interns making your product better.

Corn

Wait, wait, wait. Unpaid interns? That sounds a bit exploitative when you put it that way. But I guess if they are doing it voluntarily because they want to build cool stuff, it works out. But what about the safety side? We keep hearing that these models are dangerous. If you give the weights to everyone, can't the bad guys just take the safety filters off?

Herman

That is the core of the disagreement between the big labs. Companies like OpenAI argue that open sourcing powerful models is like giving everyone the instructions to build a bioweapon. They say we need to keep the models behind an application programming interface, or API, so they can monitor how people are using them.

Corn

And I actually kind of agree with them there. I mean, I'm a cautious guy. If there is even a small chance that an AI could help someone do something truly terrible, shouldn't we keep the lid on it?

Herman

See, I actually see it differently. I think the security through obscurity argument is a bit of a fallacy. If only one or two companies have the most powerful AI, that is a huge central point of failure. If those companies get hacked, or if they decide to turn evil, we are in trouble. But if the technology is decentralized, we have a whole ecosystem of people building defenses. It is the same argument as open source software in general. Linux is more secure because anyone can see the code and find the holes.

Corn

But software code is different from a giant pile of numbers that represents a neural network. You can't just read the weights and see where the evil is. It is a black box. So if I take a model and I fine tune it to be a jerk, how does having it be open source help anyone stop me?

Herman

Because it also allows the researchers to build better detection tools. You can't study the biology of a virus if the virus is locked in a vault you aren't allowed to enter. By having the models open, we can understand their limitations and their biases much better.

Corn

I'm still not entirely convinced, but I see the logic. Let's pivot for a second. We've talked about Meta, but what about the smaller players? Like Mistral in France. They are a relatively small team compared to the giants in Silicon Valley, yet they are putting out models that are beating some of the closed ones. Why are they doing it?

Herman

Mistral is a fascinating case. They are positioning themselves as the European alternative to the American giants. By being open, they are gaining massive adoption very quickly. It is a way to bypass the massive marketing budgets of Google or Microsoft. If your model is good and it is open, the community will do the marketing for you.

Corn

It's like the indie band that gives their first album away for free on the internet to build a following before they start charging for concert tickets.

Herman

Exactly. And speaking of things that aren't free, we should probably hear from the guy who keeps our lights on.

Corn

Oh, right. Let's take a quick break for our sponsors. Larry: Are you tired of your garden looking like a boring green rectangle? Do you wish your petunias had more... personality? Introducing Glow-Gro Bio-Luminescent Dust! Just sprinkle a handful of our proprietary, non-regulated mineral powder over your lawn, and watch as your grass emits a haunting, neon purple glow that can be seen from low earth orbit. It is perfect for late-night barbeques or signaling to long-lost civilizations. Is it safe for the soil? We prefer the term "transformative." Will it make your neighbor's dog bark at nothing for hours? Probably! But beauty has a price, and that price is very reasonable. Glow-Gro - make your backyard the only thing visible during a power outage. BUY NOW!

Herman

...Alright, thanks Larry. I am not sure I want my lawn visible from space, but I appreciate the enthusiasm.

Corn

I think my neighbor already uses that stuff. His bushes have been humming lately. Anyway, back to AI. We were talking about why companies give this stuff away. We've got developer mindshare, talent acquisition, and community optimization. But is there a world where open source eventually wins and the closed models just... disappear?

Herman

That is the trillion dollar question. There is a leaked memo from a Google engineer from last year titled "We Have No Moat, and Neither Does OpenAI." The argument was that while Google and OpenAI were busy competing with each other, the open source community was quietly eating their lunch. They were making models run on laptops that used to require a whole server room.

Corn

I remember that! It was after that model from Meta, the first Llama, got leaked, right?

Herman

Right. It wasn't even supposed to be fully open, but once the weights were out there, the community went wild. Within weeks, people had it running on iPhones and Raspberry Pis. That is the power of the crowd. When you have ten thousand hobbyists working on a problem because they think it is fun, they can often move faster than a thousand engineers at a big corporation who have to sit through meetings all day.

Corn

But there is a limit, right? You still need those massive clusters of high-end chips to train the base model. A hobbyist can't do that in their garage. So doesn't that mean the big guys still hold all the cards?

Herman

To some extent, yes. The initial training is the barrier to entry. But once that base model exists, the open source community can take it and specialize it. They can make a version that is amazing at writing poetry, or a version that is a master at coding in a specific language. The closed companies have to try to make one model that is good at everything, which is much harder.

Corn

Okay, so let me take the other side here. If I am an investor in OpenAI, and I see Meta giving away Llama for free, I am terrified. Because how do I charge for my model if there is a free version that is ninety-five percent as good?

Herman

Well, that is why OpenAI is pivoting so hard toward the "agent" model and deep integration with Microsoft. They are betting that being "good enough" isn't enough. They want to be the best, and they want to be everywhere. They are selling the reliability, the support, and the early access to the absolute cutting edge. But you're right, it puts a massive downward pressure on prices. It is great for us as consumers, but it is a brutal war for the companies.

Corn

It feels a lot like the early days of the internet. Remember when people thought you could charge for an email account? And then Gmail came along and said, "it is free, just give us your data." Except here, Meta isn't even asking for our data in the same way. They are just giving the model away.

Herman

True, but they are protecting their own ecosystem. If AI becomes a commodity—something that is cheap and available everywhere—then the value shifts to the platforms where people actually use the AI. And who owns the biggest platforms? Meta. If AI is free, people spend more time on Instagram and WhatsApp using AI features, and Meta wins anyway.

Corn

I don't know, Herman. That feels like a bit of a stretch. It's like saying a car company would give away engines for free just so people would spend more money on car fresheners. The engine is the expensive part!

Herman

But if giving away the engine means that everyone starts building their cars to fit your specific engine, then you control the standards of the entire automotive industry. You become the infrastructure. That is a very powerful place to be.

Corn

Hmm. I guess. It still feels like a massive gamble. Speaking of people who have strong opinions on gambles, I think we have someone on the line.

Herman

Oh boy. Is it...?

Corn

Yep. It's Jim from Ohio. Hey Jim, you're on My Weird Prompts. What do you think about all this open source AI stuff? Jim: Yeah, this is Jim from Ohio. I've been listening to you two eggheads talk about "ecosystems" and "mindshare" for the last fifteen minutes and I gotta tell ya, you're missing the forest for the trees. It's a tax write-off! That's all it is! My neighbor Jerry tried to tell me it's about "the future of humanity," but Jerry also tried to tell me that his lawnmower can run on vegetable oil. It can't. The whole street smelled like french fries for a week and his grass is still six inches high.

Corn

A tax write-off, Jim? That seems like a lot of effort just to save some money on the back end. Jim: Everything is a tax write-off if you're big enough! But that's not even my main point. You guys are acting like this is some big gift to the world. It's just more junk! In my day, if you wanted to know something, you went to the library and you looked it up in a book that didn't hallucinate or tell you to put glue on your pizza. Now we've got these "open weights" which sounds like something I should be doing at the gym for my bad hip. By the way, the weather here is miserable. It's been raining sideways for three days and my cat Whiskers refuses to go near the window. He just sits there and glares at me like it's my fault.

Herman

I'm sorry about the rain, Jim, and Whiskers' attitude. But don't you think there's value in having these tools available to everyone, rather than just locked up in a few big companies in Silicon Valley? Jim: Available to everyone? You mean available to the scammers! I already get ten calls a day from "Microsoft" telling me my computer has a virus. Now you're telling me they're going to have free AI to make those calls sound like my own sister? No thank you! Keep it locked up! Put it in a vault and throw the key in the lake! You guys are too young to remember when things were simple. Now if you'll excuse me, I have to go see if Jerry is trying to put salad dressing in his leaf blower.

Corn

Thanks for the call, Jim. Always a pleasure.

Herman

He's not entirely wrong about the scamming part, though. That is a real risk. When you open source these models, you are giving the tools to everyone, including the people who want to generate deepfakes or phishing emails.

Corn

But Jim's point about the library is what got me. Is there a risk that by making AI so easy and free, we stop doing the hard work of actually learning things? If I can just ask a free Llama model to summarize a book for me, am I ever going to read the book?

Herman

That is a broader philosophical question, Corn. But I would argue that open source actually helps there. If the models were only held by a few companies, they could decide what information you get to see. They could sanitize history or bias their answers toward their own corporate interests. With open source, you can have a model that is trained to be as objective as possible, or one that represents a specific cultural viewpoint. It prevents a monopoly on truth.

Corn

I like that. A "democratization of intelligence," as the tech bros like to say. But let's get practical. If I'm listening to this, and I'm not a computer scientist, does open source AI actually affect my life right now?

Herman

Absolutely. It's the reason why your phone's photo editor is getting so much better. It's the reason why there are suddenly a thousand new startups building everything from AI lawyers to AI doctors. Most of those startups couldn't afford to pay OpenAI millions of dollars a year in fees. They are building on top of open models. It lowers the barrier to innovation. It means the next big thing might come from a kid in a bedroom in Nairobi instead of just a guy in a hoodie in Palo Alto.

Corn

Okay, that is a pretty cool vision. But what about the future? Do you think the "closed" models will eventually lose their lead? Right now, GPT-4 and Claude 3.5 are still arguably better than the best open source models. Will the gap ever close?

Herman

I think the gap is shrinking faster than people realize. Every time a new closed model comes out, the open source community catches up in a matter of months. And there is a theory called "distillation." You use the really big, expensive, closed model to help train a smaller, open model. It's like a teacher showing a student how to solve a problem. The student learns much faster than if they had to figure it out from scratch.

Corn

Wait, isn't that against the rules? I thought OpenAI had terms of service that said you couldn't use their output to train other models.

Herman

They do. But how do you prove it? And in the open source world, people are doing it anyway. It's the "Wild West" of data. This brings up another point where I think you and I might disagree. I think the copyright issues around AI training are eventually going to favor open source.

Corn

Why? I would think the big companies with the big legal teams would be better at handling copyright.

Herman

Nah, because the big companies are "sue-able." You can take Google to court. You can't really sue a decentralized collection of thousands of anonymous developers who are contributing to an open source project. If the data is out there, and the model is out there, you can't put the toothpaste back in the tube.

Corn

I don't know, Herman. That sounds a bit like you're advocating for digital anarchy. I think we need some kind of framework where creators get paid for their work, even if it's used to train an AI. Just because it's open source doesn't mean it should be a free-for-all with other people's intellectual property.

Herman

I'm not saying it's right, I'm just saying it's the reality of how the technology is moving. The friction of the law usually loses to the fluidity of the code. But look, we should talk about the "so what" for the listeners. What are the takeaways here?

Corn

Well, for me, the first one is that "free" doesn't mean "charity." When you see a company like Meta giving away a powerful tool, look for the hidden strategy. Are they trying to set a standard? Are they trying to hurt a competitor? There is always a reason.

Herman

Exactly. And the second takeaway is that the "open versus closed" debate is going to define the next decade of technology. It's not just about AI; it's about who controls the infrastructure of our lives. If you care about privacy and autonomy, you should probably be rooting for the open source side, even if it's a bit messier.

Corn

And I'd add a third one: don't count out the hobbyists. Some of the most interesting stuff in AI right now isn't happening in boardrooms; it's happening on Discord servers and GitHub repositories where people are just messing around and seeing what happens.

Herman

I agree with that. The sheer speed of innovation in the open community is breathtaking. It's a great time to be a nerd, Corn.

Corn

It's a confusing time to be a sloth, Herman. But I think I understand it a little better now. It's about power, standards, and a whole lot of expensive computer chips.

Herman

That's a pretty good summary. And maybe a little bit of Larry's glowing dust.

Corn

Let's hope not. My eyes are sensitive enough as it is. Well, I think that covers it for today's dive into the world of open source AI. Big thanks to Daniel for sending in that prompt—it really got us thinking about the house we live in, figuratively and literally.

Herman

Mostly literally. I still think he needs to do something about that squeaky floorboard in the hallway.

Corn

One thing at a time, Herman. If you enjoyed this episode, you can find My Weird Prompts on Spotify, or head over to our website at myweirdprompts.com. We've got an RSS feed there for the subscribers and a contact form if you want to send us your own weird prompts. We are also on all the other major podcast platforms.

Herman

And thanks to Jim from Ohio for calling in, even if he thinks we're just a tax write-off. Say hi to Whiskers for us, Jim!

Corn

Until next time, I'm Corn.

Herman

And I'm Herman Poppleberry.

Corn

Keep it weird, everybody. Bye!

Herman

Goodbye