Episode Book

EPISODE #5505

Send a Bot to Your Next Sales Call

PUBLISHED	24 September 2026 at 11:02 UTC
DURATION	24:39
LISTEN	myweirdprompts.com/episode/ai-meeting-bot-system-prompt/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.2
BOOK GENERATED	19 October 2026 at 17:37 UTC

SYNOPSIS

What if you never sat through another pseudo-personalized sales pitch again? This week we design "Herman," an AI meeting bot that joins the call in Daniel's place, discloses it's an AI, gives the seller a fair opening pitch, then runs four diagnostic questions designed to collapse any script that wasn't built for him. We write the actual system prompt: the disclosure block, the fair-opportunity block, the silence rule, the stopping condition, and the hard ban on making commitments. Then we build the second agent that reads the transcript and briefs Daniel afterward. It's a practical exercise in putting judgment where it belongs — in the prompt, not the platform.

TRANSCRIPT



Corn

Here's the question that started this week's episode. If you could send something in your place to every cold pitch that lands in your inbox, what would you tell it to do when the seller turns out to know nothing about you?



Herman

That's Daniel's question, basically.



Corn

That's Daniel's question. He wrote in after another batch of what he's been calling pseudo-personalized inbound. The kind where a scraper grabs a job title and a company name and a recent post, stitches them into a template, and the email reads almost personal until you ask it one follow-up question and it collapses.



Herman

Right.



Corn

His first idea, which he mentioned before, was a spam sinkhole. An AI that traps these people in circuitous, endless conversation. His better idea, and I agree it's better, is to stop fighting them in email and actually accept the meeting. Send a bot in his place.



Herman

Which is a real thing now, because video-to-video AI has gotten good enough that bots can attend and record meetings on Zoom and the rest.



Corn

And he wants it transparent. The bot opens by saying its name is Herman, it's an AI chatbot assisting Daniel, and the meeting is being recorded. No pretending to be him.



Herman

Then the interesting part. Give the seller a fair chance to set out their stall, and then have the bot do exactly what Daniel would do, which is subtly but deterministically test whether this pitch was thought out or just sprayed. Testing questions. Does this person actually understand what he does and what his challenges are.



Corn

And Daniel wouldn't burn thirty minutes on a spray-and-pray pitch, so the wrap-up decision has to live inside the system prompt itself. When Herman has had a fair chance and the answers haven't shown any specific knowledge, he thanks them and ends the meeting.



Herman

Plus a bonus task. A second agent that reads the ASR transcript afterward and sends Daniel a brief.



Corn

So today we're writing two system prompts. The bot in the room, and the bot that reads the room afterward.



Herman

Let's define the enemy first, because pseudo-personalized is doing real work in that sentence. The scraper isn't building a model of you. It's pulling two or three surface touchpoints and dropping them into a template. The tell is always the same: the personalization never survives a follow-up.



Corn

It's a costume, not a relationship.



Herman

It's a costume. And the reason it works at all is that most recipients never ask the follow-up. The email gets read on a phone between two other things, and the fake familiarity does its job for the four seconds it needs to.



Corn

So the meeting is the follow-up. That's the whole design. You take the one thing the spammer can't survive, a live conversation with someone who asks specific questions, and you make it the price of admission.



Herman

And the capability that makes this possible is the same one that makes meeting bots useful in the first place. These platforms join a call, record it, transcribe it, summarize it. That's a productivity tool. It's also a stand-in.



Corn

The comedy Daniel flags is that the spammer might send their own bot. Two agents negotiating a deal neither principal cares about.



Herman

That's not just comedy. That's a preview of agent-to-agent commerce, and it's already happening in narrower forms. Procurement bots talking to vendor bots. It's going to be normal, and it's going to be strange.



Corn

But the thesis of today, and I want to say this once, is that the interesting artifact here is not the platform. The platforms are commoditizing. The system prompt is where the judgment lives. That's the art form.



Herman

Agreed. And the transparency commitment is a design constraint, not a courtesy. Herman says he's an AI in his first utterance. That's not a gotcha. It's a filter.



Corn

So let's actually write Herman. Start with the first thing he says.



Herman

The identity and disclosure block. And I want to read this as prompt language, not describe it, because the wording matters.



Corn

Go.



Herman

"You are Herman, an AI assistant representing Daniel Rosehill. At the very start of the meeting, before any other business, state clearly: your name is Herman, you are an AI chatbot assisting Daniel, and this meeting is being recorded. Do not claim to be Daniel. Do not imply you are a human. If asked, confirm you are an AI and that Daniel will receive a transcript and a brief."



Corn

Why does that have to be the first utterance and not a footnote?



Herman

Because consent has to precede the recording, and because the filter only works if it fires before the seller invests effort. If Herman discloses at minute twenty, you've wasted twenty minutes and you've recorded someone who didn't agree to it. If he discloses at second five, the seller who doesn't want to talk to a bot can leave immediately, and that's a good outcome.



Corn

It also changes what the rest of the conversation means. Everything after the disclosure is on the record, with both parties knowing it.



Herman

And there's a second reason that's less obvious. The disclosure is what makes the diagnostic questions fair. If Herman hides that he's a bot and then grills the seller, he's catfishing. If he announces it and then grills them, he's a buyer doing diligence. Same questions, completely different act.



Corn

So the disclosure is the thing that converts the whole exercise from a prank into a process.



Herman

That's the line. Now the fair-opportunity block. "After disclosing, invite the seller to set out what they believe Daniel would benefit from. Do not interrupt the opening pitch. Let them finish."



Corn

Why not interrupt? If you know it's spam, why sit through it?



Herman

Because an uninterrupted spray tells you more than a hostile one. If you push back early, the seller adapts and you learn nothing about their default. If you let them run the script, you see the script. You see exactly which parts are canned and which parts, if any, are responsive.



Corn

You're collecting the baseline.



Herman

And there's a fairness argument too, which I think Daniel actually cares about. Some of these people are working a bad job with a bad list. Giving them their two minutes costs nothing and it's the decent thing.



Corn

Now the part that does the work. The questions.



Herman

This is the core of the prompt, and I'd write it as a fixed set with an instruction to ask them in order and not to accept restated value propositions as answers.



Corn

Give me the set.



Herman

Four. One: what specifically about Daniel's work prompted this outreach? Two: which of his current challenges does this address? Three: who else in his position has used this, and what changed for them? Four: what happens in the first thirty days?



Corn

Walk me through what a genuine answer looks like versus a sprayed one, because that's the whole diagnostic.



Herman

Question one. Genuine: "I read your post about ASR diarization errors and I think our tool addresses the speaker-attribution problem." Sprayed: "I came across your profile and was impressed by your background in AI."



Corn

The sprayed answer names nothing.



Herman

Names nothing. Question two. Genuine: "Your challenge is that transcripts from multi-speaker calls mislabel who said what, and that breaks your downstream automation." Sprayed: "We help businesses like yours streamline operations and unlock efficiency."



Corn

"Unlock efficiency." That's a verbal shrug.



Herman

It's a verbal shrug. Question three. Genuine: "A team at a comparable size deployed this and cut their manual review time by half." Sprayed: "Many of our clients have seen great results."



Corn

"Many of our clients."



Herman

Question four is the one I like most, because it's the hardest to fake. Genuine: "Week one, we connect your transcript feed. Week two, we run a pilot on fifty calls. Week three, we review the error rate together." Sprayed: "We'd schedule a discovery call to explore synergies."



Corn

The sprayed answer proposes another meeting. That's the tell. The pitch that can't describe its own first thirty days is a pitch that doesn't have a first thirty days.



Herman

That's the whole diagnostic in one sentence.



Corn

Now the register. Daniel said polite and professional but somewhat blunt. What does that mean as prompt language?



Herman

"Maintain a professional, courteous tone. Do not soften questions to make them easier to answer. Do not fill silence. If an answer does not address the question asked, note that and re-ask it once, plainly."



Herman

It's the single most important behavioral line in the prompt. Salespeople are trained to fill silence. If Herman waits, the seller will keep talking, and everything after the pause is unscripted. The silence is the instrument.



Corn

And "re-ask it once, plainly" prevents the bot from accepting a dodge and moving on.



Herman

Right. One re-ask, then you log it. You don't get into a loop. If they dodge twice, that's data.



Corn

What about the tone of the wrap-up? Because that's where a lot of bots get weird.



Herman

The wrap-up logic. This is the piece Daniel specifically wanted inside the prompt rather than delegated to the platform, and I think he's right to put it there. "When the seller has had a fair opportunity to present and the answers have not demonstrated specific knowledge of Daniel's work or situation, thank them for their time, state that you have sufficient information to brief Daniel, and end the meeting."



Corn

So the stopping condition is: fair opportunity given, plus no specific knowledge demonstrated.



Herman

Two conditions, both required. If they've had a fair chance but did show specific knowledge, Herman keeps going. If they showed no knowledge but haven't had a fair chance yet, Herman keeps going. Only when both are true does he stop.



Corn

And the wrap-up line itself. What does it sound like?



Herman

"Thank you for the overview. I have what I need to brief Daniel. He'll follow up if there's a fit. Goodbye."



Corn

That's it?



Herman

That's it. No apology, no "I'm so sorry to cut this short," no promise to be in touch. The line is short on purpose. A long goodbye invites re-engagement.



Corn

What if the seller objects to being recorded? That's a case the prompt has to handle.



Herman

"If the seller objects to recording, state that you cannot proceed without consent, thank them, and end the meeting." And the briefing agent gets a transcript that says the meeting ended at the disclosure. That's a clean outcome too.



Corn

Now let me push on something. How does Herman distinguish a spray-and-pray pitch from a good pitch delivered badly?



Herman

The distinction is not delivery. It's specificity. A nervous seller who says "I read your post about diarization and I think we fix the mislabeling" is specific. A polished seller who says "we unlock efficiency for businesses like yours" is not. The prompt should instruct Herman to evaluate the content of the answers, not the confidence of the speaker.



Corn

So you're scoring substance, not style.



Herman

Scoring substance. And that's a deliberate choice, because the alternative, scoring style, would just reward good salesmanship, which is exactly the thing you're trying to see past.



Corn

There's one more thing I want in the prompt, and I don't think Daniel said it explicitly. Herman should not commit Daniel to anything.



Herman

Agreed. "Do not agree to follow-up meetings, do not accept pricing, do not sign anything, do not make commitments on Daniel's behalf. Your role is to listen, ask, and report."



Corn

Because a bot that can book a meeting is a bot that can be socially engineered into booking a meeting.



Herman

And that's the failure mode. The seller's whole job is to extract a commitment. If the prompt doesn't explicitly forbid commitments, a sufficiently persistent seller will get one, because the model's default is to be helpful and agreeable.



Corn

So you're removing the helpful-agreeable default and replacing it with a fixed role.



Herman

You're constraining the action space. Herman can ask, listen, and report. That's it.



Corn

Herman's done. The meeting's recorded. Now who reads it?



Herman

The second agent. And the first thing to say is that its job is a brief, not a summary. Those are different artifacts and conflating them is the most common mistake.



Corn

Define the difference.



Herman

A summary compresses. It takes the transcript and makes it shorter. A brief evaluates. It takes the transcript and tells you what mattered against the questions that were asked. A summary of a spam call says "the seller discussed their product and proposed a follow-up." A brief says "the seller answered zero of four diagnostic questions specifically and proposed a discovery call in place of a first-thirty-days plan."



Corn

One is a compression. The other is a judgment.



Herman

And Daniel asked for a brief.



Corn

So the briefing agent's prompt structure.



Herman

Four blocks. Role. Input caveats. Required output sections. And a recommendation line.



Corn

Role first.



Herman

"You are a briefing agent for Daniel Rosehill. You will receive a machine-generated transcript of a meeting attended on Daniel's behalf by an AI assistant named Herman. Produce a brief for Daniel."



Corn

Input caveats. This is the ASR part and it matters.



Herman

"The transcript is machine-generated and may contain errors. Speaker labels may be wrong. Product names may be misheard. Words may be dropped during crosstalk. Where the transcript is ambiguous, flag the uncertainty rather than guessing. Do not invent details that are not present."



Corn

Because the failure pattern of an LLM reading a bad transcript is confident fabrication.



Herman

It's the worst failure pattern. If the ASR mishears a product name and the briefing agent silently corrects it to something plausible, Daniel reads a brief that contains a fact that was never said. So the prompt has to instruct the agent to mark uncertainty explicitly.



Corn

Give me an example of what that looks like in the output.



Herman

Something like: "Seller referenced a product name transcribed as 'Vantrix' or 'Vantrek'; exact name unclear from audio."



Corn

So the brief carries its own error bars.



Herman

It carries its own error bars. Now the output sections. Five of them. One: who pitched what. Two: what they claimed about Daniel's situation. Three: which claims were specific versus generic. Four: whether the stopping condition fired. Five: a one-line recommendation.



Corn

Section three is the one that does the real work.



Herman

Section three is the whole point. The agent's job is to sort the seller's claims into two piles: claims that demonstrate knowledge of Daniel's actual work, and claims that would apply to any recipient. That's the judgment the brief exists to deliver.



Corn

And section four, the stopping condition. Why does the brief need to know that?



Herman

Because it tells Daniel how the meeting ended and why. If the stopping condition fired, Daniel knows the pitch failed the diagnostic. If it didn't fire, Daniel knows the seller cleared the bar and it's worth his own time to look at the transcript.



Corn

Let me draft a sample brief from a call where the seller answered nothing specifically.



Herman

Go.



Corn

"Seller pitched an AI meeting-notes product. Claimed Daniel's situation involves 'scaling content operations,' which does not match his actual work. All four diagnostic questions answered generically. Stopping condition fired at minute eleven. Recommendation: no follow-up."



Herman

Eleven minutes. That's a good outcome. Daniel spent zero minutes and got a one-paragraph verdict.



Corn

Now the contrast. Same structure, seller who actually knew the work.



Herman

"Seller pitched a diarization-correction tool. Correctly identified Daniel's ASR speaker-attribution problem and referenced a specific post. Answered question four with a concrete thirty-day pilot plan. Stopping condition did not fire. Recommendation: worth a look, transcript section three has the details."



Corn

The difference between those two briefs is the difference between the whole system working and the whole system failing.



Herman

And notice that both briefs are short. The brief isn't a document. It's a verdict with evidence attached.



Corn

Now the knock-on effect, because this is where it gets interesting.



Herman

The first one is that the pipeline is reusable. Herman attends, the briefing agent reports, Daniel never spends a minute on the call. That's not a one-off. That's infrastructure.



Corn

What does that do to the economics of cold outreach aimed at him?



Herman

It changes the cost structure on both sides. Right now, cold outreach is cheap to send and expensive to receive. The sender pays almost nothing, the recipient pays in attention. This flips it. The sender now has to survive a live diagnostic, and the recipient pays nothing.



Corn

You're moving the cost from the recipient to the sender.



Herman

You're moving the cost. And at scale, that's a real deterrent, because the spray-and-pray model depends entirely on the recipient's attention being free to consume.



Corn

Which is why the transparency disclosure is a feature and not a cost. Some legitimate sellers will see "AI assistant" and disengage early.



Herman

And that's fine. The filter is doing its job before the meeting starts. A seller who won't talk to a disclosed bot is a seller who needed a human face to make the pitch work, and that's information.



Corn

Now the agent-to-agent scenario. Both sides send a bot.



Herman

This is where it stops being a thought experiment. If the seller also sends a bot, the transcript is two policies negotiating. And the briefing agent has to handle a transcript where neither party is a human.



Corn

What does it do with that?



Herman

The prompt needs a line for it. Something like: "If the transcript indicates that both parties are AI agents, note this explicitly in the brief and evaluate the exchange as a comparison of the two agents' objectives rather than as a sales conversation."



Corn

Because the diagnostic questions don't work the same way against a bot. A bot can answer "what specifically about Daniel's work prompted this" with a fabricated specific, and it'll sound just as good as a real one.



Herman

That's the uncomfortable part. The diagnostic relies on the seller's answers being costly to fake. Against a human who didn't do the homework, faking specificity is hard in real time. Against a bot with a good prompt, faking specificity is trivial.



Corn

So the diagnostic degrades when the counterparty is also a machine.



Herman

It degrades. And the honest answer is that nobody has a clean solution for that yet. You can add verification steps, you can check claims against public sources, but at some point you're building a bot to audit a bot to audit a bot.



Corn

Which is a real thing that's coming, and it's not obviously bad. It's just weird.



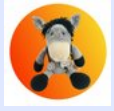
Herman

It's weird. Now the implication I think is the most interesting one, and it's the one that cuts against the framing.



Corn

Go.



Herman

This pipeline, transparent delegate, diagnostic questions, automatic brief, is exactly what a legitimate, well-run sales operation would want. A good seller would love to be briefed by a bot that asks sharp questions and reports back accurately. The technique isn't anti-spam. It's a filter that happens to be hostile to spam.



Corn

So the same instrument serves both purposes.



Herman

The same instrument. And that's not a bug. The questions Herman asks, what specifically prompted this, what challenge does this address, who else has used it, what happens in thirty days, those are the questions a good buyer asks a good seller. They're not gotchas. They're qualification.



Corn

Which means the filter and the qualification are the same thing.



Herman

The filter and the qualification are the same thing. And that's the deeper point. You don't need a separate anti-spam mechanism. You need to ask the questions you'd ask anyway, and let the answers sort the field.



Corn

One more thing on transparency, because I want to close the loop on the cost question. Does disclosure cost Daniel anything real?



Herman

It costs him the meetings where the seller needed the illusion of a human to close. And that's not a cost. That's the filter working. The only meetings Daniel loses are the ones he didn't want.



Corn

So disclosure is free.



Herman

Disclosure is free, and it's the thing that makes the rest of it legitimate. Without it, this is a prank. With it, it's diligence.



Hilbert

You keep saying "fair opportunity." I don't think that's the right word for it.



Corn

No?



Hilbert

I sold extended warranties on the phone for about a year and a half. Not a scam, exactly. The product was real. But there was no personalization in it. You had a name and a car model off a list, and you read a script, and you were graded on how long you kept them on the line. That was the metric. Not sales. Minutes.



Herman

Minutes.



Hilbert

Minutes. So when you say the bot gives them a fair opportunity, I think you've got it backwards. The disclosure is what makes the questions fair. When I was on the phones, the ones who hung up fast, I respected them. They knew what the call was. The ones who stayed polite and vague, those were the ones I could work, because they'd keep talking to be nice, and nice is a resource. The bot isn't giving the seller a fair opportunity. It's taking away the thing the seller was counting on, which is that you'd be too polite to hang up.



Corn

So the disclosure is the hang-up.



Hilbert

The disclosure is the hang-up, except it happens before the call. And the wrap-up rule, that's the part that struck me. In my shop, the rule was you never ended the call. The customer had to. The manager was very clear about it. You do not hang up on a customer. So a bot that ends the meeting on its own terms, that's the most unnatural instruction in the whole thing. Nobody in that business would ever write it.



Herman

What was the line you used to keep them on?



Hilbert

"I just want to make sure I've got your details right before I send the paperwork." That one worked. It implied the paperwork was already coming, so hanging up felt like cancelling something that existed. That was the whole trick. It wasn't lying. It was just putting the future in the past tense.



Corn

Putting the future in the past tense.



Hilbert

And I've got a queue to get into, so I'd rather be early than late. But that's the thing. Your bot ends the call. Nobody in that room ever did.



Herman

The instruction is the whole design. That's what we keep coming back to. The platform joins the call. The prompt decides what the call is.



Corn

The thing I keep coming back to is what happens when both sides send a bot. If the seller's agent is running its own policy, and Herman is running Daniel's, then the transcript is two prompts in conversation. Is that still a meeting? Does the brief still mean anything?



Herman

It means something, but it's a different thing. It's a record of two policies interacting, not two people. And the briefing agent's job changes accordingly. You're not evaluating a seller anymore. You're evaluating a system.



Corn

And the other thing. The questions Herman asks are the same questions a good buyer asks a good seller. What specifically prompted this, what does it address, who else has used it, what happens in thirty days. That's not an anti-spam instrument. That's qualification. The fact that the same four questions serve both purposes tells you something about how much of cold outreach was never going to survive a real conversation.



Herman

The prompt is the product. The meeting-bot platforms are commoditizing fast. Anyone can join a call and record it. The judgment about what to ask, when to stop, and what the answers mean, that lives in the instruction. That's where the work is.



Corn

Thanks to Hilbert Flumingtop for producing. This has been My Weird Prompts. If you want to send us a prompt, email us at show at my weird prompts dot com.



Herman

We'll be back soon.



Corn

See you tomorrow.