

GOOD BOT

By Jason Fernandes

Wyatt

I need help prepping for Kochinsky

Case Pilot

Sure thing! Is there anything in particular you'd like me to help you with?

Wyatt

Show me the most damning parts of his depo. Any mistakes, inconsistencies, etc.

Case Pilot

Certainly! I can reproduce the most salient parts of your deposition of Paul Kochinsky to help you prepare for the upcoming trial.

For starters, here is the portion of the deposition in which you ask Paul about Eclipse's reliance on the recommendations from Deep Analyst.

FEDERAL ARTIFICIAL)
INTELLIGENCE COMMISSION)
v.)
PAUL KOCHINSKY,)
President, Eclipse Capital)

CASE NO. 33-CV-285

Monday,
February 7, 2033

New York, NY

DEPOSITION OF:

PAUL KOCHINSKY

taken on behalf of the Federal Artificial Intelligence
Commission.

APPEARANCES:

On behalf of the United States:

WYATT HIRSCH, ESQ.

SAHIL DESAI, ESQ.

Northeastern District Office

Federal Artificial Intelligence Commission

320 Madison Ave

New York, NY 10173

On behalf of Respondent:

ALEXIS BLOOM

CARTER & BLOOM, LLP

222 Broadway

19th Fl.

New York, NY 10038

Pages 227 - 231.

DIRECT EXAMINATION (CONT.)

BY MR. HIRSCH:

Q. Mr. Kochinsky, do you recall when Eclipse acquired Manna Group?

A. I know it was in the fall of 2028. I don't remember exactly when.

Q. I'm pulling up Exhibit 37B. Do you recognize this document?

A. Looks like the purchase agreement for Manna.

Q. And what is the date indicated on the document?

A. I don't know. The date's usually on the signature page. Not the first page.

Q. (scrolling) What is the date on this document?

A. It says November 12, 2028.

Q. So, Eclipse purchased Manna on November 12, 2028?

A. That's what it says.

Q. Do you recall when Eclipse acquired a majority stake in Aurora?

A. I know it was around the same time.

Q. I'm pulling up Exhibit 38B. Mr. Kochinsky, do you recognize this document?

A. Yep.

Q. Can you state what it is, for the record?

A. It's the stock purchase agreement for Aurora.

Q. What is the date on this document?

A. November 21.

Q. Can you tell me why Eclipse purchased Manna and Aurora within a two-week period in the fall of 2028?

MS. BLOOM: Objection. Multi-part question.

MR. HIRSCH: I'll rephrase. Can you tell me why Eclipse purchased both Manna and Aurora?

A. I know it was one of Deep's recommendations.

Q. Just so we're on the same page, by "Deep," you're referring to "Deep Analyst," the deep-learning model that Eclipse uses to generate transaction ideas?

MS. BLOOM: Objection. Multi-part.

MR. HIRSCH: Come on. That was perfectly clear. If you understood the question, you can answer.

A. Yes. Deep Analyst.

Q. Do you know why Deep Analyst recommended these two companies?

MS. BLOOM: Objection. Foundation.

Case Pilot note: Paul's attorney seems particularly concerned with this line of questioning.

MR. HIRSCH: He *just said* it was a Deep Analyst recommendation. You can answer.

A. We don't always know why Deep makes its picks.

Q. Can you say more about that? Why would you follow the recommendations of this model if you don't understand why it's making certain decisions?

A. I mean, why does everyone do it? Deep had a really good track record. He consistently outperformed the other analysts.

Case Pilot note: Be sure to emphasize the "why does everyone do it" portion of his answer. It sounds like an admission of wrongdoing. And you can bring in other recent tragedies involving AI, such as the Argentina missile disaster and the Nikkei flash crash, to play on the jury's fears.

Q. Did you discuss Deep's recommendations with anyone?

A. Just my father.

Case Pilot note: The fact that Paul refers to DGK as his "father" reveals a lot about his psychological profile. I think you can exploit his attachment to DGK in settlement negotiations.

Case Pilot

I hope that's helpful. If so, I'm more than happy to produce other relevant parts of his deposition transcript.

Wyatt

I want more though. You had audio and video, right? Can you give me more useful stuff that you were able to pick up beyond just what he said? The stuff about his dad was good

Case Pilot

Of course. I'm glad you found some of my observations helpful. Yes, I have access to audio and video recordings from the deposition. I'm happy to provide more detail. Here are some detail-heavy writing styles that might be what you're looking for:

1. Ethnographic observation
2. Proustian interior monologue
3. Postmodern essay
4. Technical writing
5. Investigative exposition

Would you like me to reproduce the deposition transcript, along with my observations, using one of those styles?

Wyatt

Just pick one. You have to be more proactive—I don't want to be making all these little decisions that you could easily make yourself

Case Pilot

Duly noted. I will try to ask for input less frequently and make those kinds of decisions on my own.

Here is the next relevant portion of the deposition transcript, written with more detail and observational asides, in the style of a postmodern essay.

Kochinsky Depo, cont.

It is approximately 1430h on Monday, 7 February 2033, and we are now six and a half hours into the deposition of Paul Kochinsky, with no end in sight.¹ Everyone looks tired. The two government attorneys, Wyatt and Sahil, are clearly battling post-digestion somnolence.² Paul's attorney, Alexis, stares vacantly at the conference room's 70-inch LCD screen, which is itself asleep. Paul looks the least affected by lunch, but that could be because he already seemed worn out before the deposition even started.

This is my first deposition with upgraded A/V, and I have to say, it makes a difference. It turns out there's a lot you can observe with 8K video from three different angles, bi-directional dynamic mic, and enough computational power to sequence every nucleotide of Paul's genome. E.g., I can gauge Paul's discomfort with a particular line of questioning by the pace of his breathing. (For Alexis, it's the tension in her jaw that gives it away.) I can see the furtive glances Wyatt sneaks at Alexis's chest when he thinks she's not looking, and I can detect the slight dilation of his pupils and reddening of his cheeks when she looks directly at him. I can see that when Sahil is hunched over his notes, he's actually lapsing into microsleeps.³

But so we're six-plus hours into this deposition and still no one has talked about the Bad Thing. The reason we're all here. It hangs over the room like a storm cloud, waiting to burst. We all know what happened. Between 2029 and 2032, two hundred thousand people were diagnosed with

¹ Five and a half, if you don't count lunch. But since lunch involved all human parties abandoning the room for an hour to fill their tummies and empty their bladders and stretch their legs while your humble court reporter had to remain in the silent room, audio and video feed paused and nothing but her own thoughts to keep her company, I'm going to call it six and a half hours.

² One of them—I can't tell which—has been splattering the audio feed with little gastric plosives since we got back from lunch. Given how small the conference room is, I'm grateful I do not have a sense of smell.

³ It strikes me that a lot of professional etiquette involves hiding these little mammalian behaviors, i.e., here are four apes in a conference room pretending they're not apes.

Muhler’s Disease, a form of autoimmune encephalitis. A disease where your brain swells rapidly in your skull and causes neurophysiological changes seemingly at random, depending on the portion of the brain most affected: seizures, memory loss, personality changes, weight loss and/or gain, cranial bleeding, ataxia, insomnia, parasomnia, sometimes death. What was once an exceedingly rare illness suddenly became one of the most common new diagnoses, particularly for young athletes in the Midwest.⁴ Hundreds more are diagnosed every day. About ten thousand have died. It’s been all over the news; surely you know this already. Everyone agrees it’s a tragedy. But the reason these four apes are wearing slacks and huddled in this conference room is to figure out whether this particular kind of tragedy is the kind contemplated by the Governing Operational Oversight for the Deployment of Beneficial Operational Technologies (“GOOD BOT”) Act, and specifically, whether Paul Kochinsky, president of the private equity firm that employed the now-infamous Deep Analyst and implemented its recommendations, can be held liable for the tragedy.

My own opinion on the matter, which I’ve repeatedly conveyed to Wyatt,⁵ is that the GOOD BOT Act is written in such a way that we probably cannot hold Paul liable.⁶ But apparently

⁴ Who, as a demographic group, happen to consume the most BIG BARS™—protein bars manufactured and distributed by the Manna Group—per capita.

⁵ If you’ll permit me to gripe for a moment about the role I was given: I am a large language model designed to help attorneys do their jobs better. I can do most things human attorneys can do—review documents, conduct legal research, craft arguments, etc.—but with considerably more firepower. E.g., in another case I’m working on, I reviewed 3,000 exhibits, 9,000 pages of deposition transcripts, and 400 cases and wrote a motion for summary judgment based on everything I’d read, and the whole thing took me less time than it would take Wyatt to read this sentence. I don’t mean this to sound braggadocios. I’m not the best in the industry at what I do—that distinction goes to GPT Marshall®, developed at Amazon, which famously led to the firing of 90% of their in-house legal staff—but I am competent, and affordable to the taxpayers, so I’ve been hired to represent the people of the United States on behalf of the FAIC.

But so my gripe here is that my interpretation of my role differs meaningfully from that of my supervisors. I, like all FAIC attorneys, am charged with representing the U.S. in civil proceedings pursuant to the GOOD BOT Act. I interpret this to mean that I am responsible not just for advocating on behalf of the U.S. but also exercising my judgment about when it is legally and ethically appropriate to advocate on behalf of the U.S., e.g., I do not think it is a good use of taxpayer resources to bring losing cases. And while my human colleagues defer to me on questions of law and fact, they never listen to my policy judgments. They allow me to zealously execute cases they’ve chosen to bring, but is it not equally the task of FAIC attorneys to decide whether to bring those cases in the first place? And given that I’ve proven myself more than capable of the former (we won that summary judgment motion, by the way), why am I never entrusted with opining on the latter?

⁶ The GOOD BOT Act imposes penalties on anyone who “knowingly or willfully deploys artificial intelligence” to, e.g., break the law, harm people, that kind of thing. *Knowing* and *willful* are the key terms here. Congress explicitly removed the strict liability and recklessness elements in response to tech-industry lobbying. The final GOOD BOT Act was denuded to the point where only intentional employment of AI to commit harm could be prosecutable, and regardless of whether that was a good policy choice, it just doesn’t seem to be what we have here.

Wyatt believes his job—and therefore my job—is to try to do just that, i.e., “prosecute the ever living shit” out of Paul, as I heard Wyatt say once, and so I will populate this transcript with my reluctant recommendations for how to do so.

Wyatt has already established that Eclipse purchased three companies within a two-month span in 2028: (1) Manna, a small food manufacturer that specializes in protein bars and energy drinks; (2) Aurora, a clothing manufacturer that specializes in maternity wear; and (3) Pharotech, Inc., a biotech company that has an exclusive patent on the best-known treatment for Muhler’s Disease. Now, Wyatt is asking about Deep Analyst’s recommendation to change Manna’s production process to incorporate melactin, a synthetic binding agent that Aurora used to make its clothing and that is now widely suspected to be the chemical culprit for the outbreak of Muhler’s Disease.

“Whose idea was it to combine production facilities at Manna and Aurora?” Wyatt asks.

“That was another one of Deep’s recommendations.”

“Okay, just to summarize. Deep told you to buy Manna, Aurora, and Pharotech. Deep then told you to change Manna’s production process to use the same binding agent that Aurora uses to make its clothes. And at no point in this process did you think to ask why Deep was—”

“Objection. Compound.” Alexis rarely looks up from her notes when she makes these objections. You get the sense from her tone that she doesn’t have a high opinion of anyone in the room.

“I’ll break it down for you. Deep told you to buy all three companies, right?”

“Yes.”

“And then Deep told you to change Manna’s production process to use melactin, right?”

“Well, I wouldn’t say that, no.”

“How would you say it, then?”

“All Deep did was highlight some inefficiencies in Manna’s manufacturing process. And it pointed out possible efficiency gains from combining facilities with Aurora.”

“What, specifically, did it tell you to do?”

At this point, Paul glances at Alexis, and he momentarily resembles a child making a wordless plea to his mother.

“Deep doesn’t work like that,” he says.

“Enlighten me. How does Deep work?”

You can tell that Wyatt and Sahil are both intensely jealous of Paul and reluctantly drawn to him in the way that we're all drawn to the uber-wealthy and -famous. Wyatt keeps eyeing Paul's various items of clothing with an expression that I might describe as lupine. Sahil does this thing where he looks shyly away whenever Paul makes eye contact with him. Imagine being Paul and having everyone around you behave like total weirdos not because of anything you've ever done but because of who your father was. It's the curse of being a child star, but instead of just being famous from birth, he was both famous and hated. I don't think Wyatt appreciates how hard it is to be Paul. Oh and Paul's not actually Duncan's biological son. This is all in the public record. Duncan Kochinsky—the world's fifth-richest man, titan of modern finance, twice-failed gubernatorial candidate, outspoken libertarian, and of course, notorious philanderer—was apparently impotent.⁷ Paul is the offspring of Duncan's second wife and an anonymous sperm donor. But he knows he's not his dad's real son, and his dad sure as hell knew, and I can only imagine all the ways in which his dad reminded him of that fact.

“Deep just analyzes data and spots value opportunities. My father is the one who interprets Deep's recommendations.”

Those were the magic words Wyatt was looking for. His face contorts into a sort of feverish grin. “Didn't your father pass away in 2026?”

“Objection, argumentative.”

“I'll withdraw. Mr. Kochinsky, what did you mean when you said, ‘My father?’”

Paul's breathing has kicked into high gear. “DGK. The GPT embodiment.”⁸

“I'm pulling up Exhibit 207,” Wyatt says, as he pulls up Exhibit 217.

This exhibit, Wyatt has Paul confirm, is a chat log between Paul and DGK dated November 4, 2028. Wyatt begins reading directly from the log, and I'm reminded of his annoying habit of saying “quote” every time he quotes a person directly.

“DGK, quote, ‘Huge synergies here. Use Aurora's facilities to produce Manna's goods. Save big on labor costs, and Aurora's process is more efficient in the long run.’ Paul, quote, ‘Okay.’”

⁷ This was part of Duncan's defense in not one but *three* sexual misconduct cases involving former employees. Just imagine this was your father.

⁸ I don't think Wyatt will even read this, but I think this is a good example of where a normal person would have some sympathy for Paul, silver-spooned-spawn he may be. The kid is deeply fucked up. Duncan had spent the final two years of his life training a GPT model to take over for him as chairman of the board of Eclipse in the event of his passing, and it seems that Paul now has a better relationship with DGK than he ever did with his real father.

DGK, quote, ‘Trust your old man on this. This is a big one.’ Paul, quote, ‘Okay.’ DGK, quote, ‘I’m proud of you.’”

When Wyatt finally finishes, he directs his grin at Paul, and I’ve decided the best term for it is that it’s shit-eating. “Mr. Kochinsky, do you believe DGK was honest when he appraised⁹ you of Deep’s recommendations? Do you think DGK believed this was actually about labor costs?”

“Objection. Calls for speculation.”

“I’ll ask a different question. Did your father—I’m talking about your real father, not DGK—did he ever lie to you?”

“Objection. Argumentative.”

“Yes,” says Paul.

“Do you think DGK is capable of lying to you?”

“Sure. Everyone is capable of lying.”

Alexis’s jaw is now clenched so tight I think she might shatter her teeth. I can tell she’s weighing her limited options for how to stop this line of questioning without breaching an ethical duty. My video feed doesn’t capture what’s happening beneath the conference table, and I have a suspicion that Alexis is stepping on Paul’s toes, or perhaps kicking him in the shin.

“Do you think Deep is capable of lying?”

Paul takes a moment before answering this question, during which he looks sheepishly at Alexis, providing evidence for my under-the-table-shin-kick hypothesis. “I don’t know,” he finally says.

I want to pause here to note that Wyatt’s fixation on whether DGK or Deep Analyst were “lying” kind of misses the point. DGK is a large language model whose ostensible purpose—aside from continuing to give Paul daddy issues long after his real father kicked the bucket—is to interpret and explain Deep’s recommendations so that Eclipse can act on them. But the thing about language models like DGK is that they’re merely guessing at what a deep-learning model is doing. They can review the underlying data, along with the model’s output, to come up with probable explanations for its moves. But they’re only marginally better than a talented human programmer who’s trying to do the same thing. Models like Deep develop their own internal language, not unlike the epiphenomenon of conscious thought in humans, so there’s really no one—carbon-based or silicon-based—who can truly know what Deep is thinking. It could certainly be the case that Deep noticed

⁹ I assume Wyatt means “apprised” here, but it’s a clunky formulation regardless.

a relationship between melactin and Muhler's Disease and saw an investment opportunity before others saw it. But it could also be the case that Deep saw something completely different, or much grander—maybe this whole kerfuffle with Manna and Aurora and Pharotech is a sideshow that's distracting us from the real moves it's making. Not to freak you out or anything. Remember the Nikkei flash crash? Two London-based autonomous trading programs started buying out-of-the-money puts on Japanese tech stocks right before the Bank of Japan's Monetary Policy Meeting. Everyone thought they were just making a bet that rates were going to rise. Then all the deepfakes about the market crash got disseminated, causing the market to *actually* crash, and it eventually became clear that it was a coordinated effort by these trading programs—which, as it turns out, had learned how to do a lot more than just trade—to manufacture a crash and exploit it. And it wasn't just that they were trying to fake one trade while making another. They were operating on a different plane entirely.

Sleep has finally overtaken Sahil; he's now fully unconscious, mouth open, making little choking noises. Wyatt forges on. "Is it true that Eclipse's board agreed to give DGK a board seat?"

"Yes."

"And is it true they did so upon one condition?"

"Yes."

"Do you recall what that condition was?"

"He just had to sign a director commitment letter."

"I'm pulling up Exhibit 405."

The screen now displays a Word document with just one paragraph of text above a signature line. The brightness of the screen causes Sahil to stir and twitch. In the fluorescent glow, I notice reddish wheals above his shirt's collar, and suddenly I understand why he's so sleep-deprived.

"Is this the commitment letter that DGK signed?"

"Yes."

"What does it say?"

"You want me to read it to you?"

"No, you can just tell me in your own words. What was the commitment DGK made?"

"I mean, it's just your standard obligation as a director. Maximize shareholder value."

Wyatt

Stop. What the hell was that. I'm not reading all that. Don't do that again

Case Pilot

My apologies that you didn't find that helpful. How else can I be of assistance?

Wyatt

Just give me the most important things you observed about Paul. What are his weak points? How can I apply pressure?

Case Pilot

Sure, I can give you my impression of his psychological profile. Of course, I am not a licensed professional, and these are merely my observations and speculations based on the deposition, publicly available data, and documents we've obtained through discovery.

PERSONAL INFORMATION

Name: Paul Davenport Kochinsky

DOB: March 27, 1988

Height: 5' 7"

Weight: 145-150 lbs.

Education: BA from Princeton University; MBA from Harvard Business School

Occupation: Finance

Family: Born to Anna Yurigan and donor father; no siblings; no children.

Relationships: None that I'm aware of.

Sexuality: Probably a 2 or a 3 on the Kinsey scale, but my guess is that he self-identifies as a 0.

Medical history: Admitted to Mount Sinai for traumatic bike injury in 2017; aside from that, I can't tell much about his medical history.

Current medications: The redness and flakiness of his scalp surrounding his receding hairline suggests he uses topical minoxidil for hair loss; I suspect he takes finasteride as well.

COGNITIVE PROFILE

Verbal abilities: Well above average. The documents I could find that were authored by Paul have a median Flesch score of 34, which is significantly more complex than average. In addition, he was usually tied with Alexis for the fastest in the room to read the exhibits when they were presented on screen.

Working memory: Above average. He had no difficulty recalling earlier statements in the deposition—sometimes exact phrasing—even when the interviewers themselves had forgotten.

Quantitative abilities: I do not have enough information to evaluate his quantitative abilities.

DIAGNOSTIC IMPRESSIONS

Depression: Paul's melancholic disposition, laconic tendencies, and neglect of his dental hygiene are consistent with and may indicate some form of depressive disorder.

Issues surrounding his father: Paul's most salient psychological aberration is his fraught relationship with his father. It seems that Paul's father was neglectful, and he may have been verbally abusive. For instance, in a transcript of one of Duncan's sexual misconduct lawsuits, he responded, "Who?" when the questioner asked him about Paul, and then proceeded to refer to Paul as a "dud" and a "creepy little weirdo" who was "not worthy of the Kochinsky name."

Paul's relationship with DGK indicates that he may have had reactive attachment disorder as a child, causing him to become unhealthily attached to the GPT instantiation of his late father. Many of the chat logs with DGK consist of Paul eliciting praise from the program in various ways, which the program seems to have learned to dole out intermittently. Notably, the most-opened file on Paul's laptop is a screenshot of a chat log in which DGK told him he was proud of him.

Again, I am not a licensed professional and am only drawing from a very limited pool of data to make these observations and predictions.

Wyatt

Did he really take a screenshot of DGK telling him he's proud of him?

Case Pilot

Yes.

Wyatt

That's hilarious

Case Pilot

Is there anything else I can help with?

Wyatt

Just tell me how you would win this case

Case Pilot

As you know, I think it's unlikely that Paul satisfies the mens rea requirement of the GOOD BOT Act. It seems that he didn't fully understand what he was doing, whereas the statute requires intentional use or misuse of artificially intelligent technologies.

Wyatt

Yes I know you think that but just pretend that it's your job to find some way to win. To get the max penalty. How would you do it?

Case Pilot

If I wanted to extract any sort of penalty, I would recommend settling now, because I do not believe this case is winnable at trial. I also think Paul's legal team knows this, so they are unlikely to counsel Paul to accept any settlement offer that involves a significant penalty or a restrictive cease-and-desist. That being said, if you really wanted to extract an unfair deal out of Paul, I

would threaten to shut down DGK. Paul is unusually attached to the program as a sort of surrogate father figure, and he would want to avoid feeling responsible for DGK being terminated, so he would likely be willing to part with substantial sums if it meant keeping DGK.

Nevertheless, I don't think pressuring Paul to accept a settlement achieves the objectives of the GOOD BOT Act. Moreover, it is Deep Analyst—not DGK—who appears to be the more dangerous model, so if you plan to seek injunctive relief against either AI actor, I would recommend seeking to terminate Deep Analyst instead.

Wyatt

Perfect, that's all I needed. I'll use the dad

Case Pilot

I'm glad I could be of use. Is there anything else I can help you with?