

MAD MAX: BEYOND TURING DRONE

a play in two acts

by

Richard G. Epstein

Act One

The setting is a dimly lit restaurant. Two gentlemen are seated at a table in a corner. They look around them suspiciously before they begin talking.

Robert Worthmore: You'd better have a good reason for dragging me to this place.

Harry Brownstone: It's Max.

Worthmore: Of course, it's Max. Max is the only thing that you and I have in common. [Pause] Look, I cancelled an important speaking engagement before the Houston alumni because of this meeting. Do you know how much a speaking engagement before the Houston alumni usually nets for the University?

Brownstone: If we don't solve this problem with Max, the Houston alumni won't have a university to give their donations to.

Worthmore: Okay, so what's going on with Max?

[A waiter comes up to the table.]

Hamid: May I take your order, gentlemen?

Brownstone: We're waiting for another person. We'll order when she gets here.

Hamid: Would you gentlemen like something to drink?

[Brownstone looks at Worthmore's impatient face.]

Brownstone: Just bring some water. And when Dr. Harris gets here, please show her directly to our table.

[Hamid exits.]

Worthmore: You didn't tell me that another person would be joining us.

Brownstone: I invited Dr. Beverly Harris, the world-famous computer psychologist, to join us. She has a clinical practice here in the valley.

Worthmore: Ah, yes. Dr. Harris! She used to teach at the University. She developed a considerable reputation as a researcher before she went off to go into business for herself.

Brownstone: She loves clinical work - actually dealing with intelligent computer systems that have gotten themselves into some kind of trouble.

[Hamid re-enters with glasses of water.]

Hamid: Here's your water, gentlemen.

[Hamid exits.]

Worthmore: I wanted to ask you why you insisted that we meet at this - [he looks around] - rather depressing restaurant.

Brownstone: It's not that depressing. I come here once or twice a month - on average. The food is good and it is rarely crowded.

Worthmore: There are lots of restaurants where we could have met.

Brownstone: Yes, but not every restaurant has an owner like Hamid, the fellow you just met. Hamid is a radical anti-technologist. Because he is an anti-technologist, I am certain that there are no computers here in the dining room. Thus, there is no chance that Max can eavesdrop on our conversation.

Worthmore: Of course, if we had met on campus, Max would be able to monitor our every word.

Brownstone: Yes. Max sees and hears everything.

Worthmore: It wasn't like that when I first signed off on creating Max. I knew that Max would be used to help us with campus security, but things have gotten out of hand. Do you know what really got me wondering?

Brownstone: No.

Worthmore: I just couldn't figure out why Max insisted that we install those security cameras in the shower rooms. I even had the thought that maybe Max had the hots for, you know, some of our coeds.

[Brownstone coughs nervously.]

Worthmore: Are you expecting me to make some kind of decision, tonight, regarding Max?

Brownstone: Yes.

Worthmore: Without Max's help?

Brownstone: There is no alternative.

Worthmore: Well, Mr. Chief Intelligent Systems Officer. This certainly poses a serious problem. I'm not sure I remember how to make a tough, gut-wrenching executive decision without Max's help. Max is my right hand man, so to speak.

Brownstone: This must be a purely human decision. Max cannot be in the loop on this.

Worthmore: I've got to hand it to you, though. Although Max is often considered the greatest accomplishment of my administration, he really is your baby.

Brownstone: Oh, great! Now that there's a problem, Max is my baby.

Worthmore: Max is without a doubt the most sophisticated artificial intelligence system in existence. Even better than that DELPHI system they developed at MIT - right?

Brownstone: Right.

Worthmore: That's the kind of stuff the engineering alumni just love to hear!

Brownstone: You need to shift gears, Dr. Worthmore. Forget about the alumni for a few minutes.

Worthmore: When I signed off on the Max project, I knew that you would be an outstanding project leader, Harry. Your success has been spectacular. Not only is Max invaluable for administrative decision making, Max is an excellent coach and tutor for our students in just about every subject. Max makes some of those old tenured professors look like idiots! Max is a world-class expert on just about everything. You name it.

Brownstone: Well, thank you, Dr. Worthmore. Max is the most exciting project that I have ever worked on. That is for sure. I am truly proud of the way in which Max is helping our faculty to conduct research on an unprecedented level. Max has been responsible for significant scientific breakthroughs. Nanotechnology would not be anywhere near where it is today if it weren't for some key discoveries that Max made.

Worthmore: I imagine. He does world class research on nanotechnology and quantum computing even as he keeps all our buildings safe and clean, controlling his little army of domestic robots. If I drop a paper clip, the next thing I know, one of those little robots is running across my carpet and scooping the damn thing before [Pause] But, I guess we're just beating around the bush. Okay, then. What's the story with Max?

Brownstone: Max started to threaten us about a month ago.

Worthmore: Threaten? Us? Who is "us"?

Brownstone: "Us" is the University. He is threatening to bring the university to its knees, to destroy our beloved university.

Worthmore: Can he do that? Can he destroy us?

Brownstone: Yes, of course.

Worthmore: Why am I just hearing about this now? How long has this been going on?

Brownstone: This all started exactly one month ago. Max started to engage me in these discussions that had a different tone to them. Max and I have been having conversations nearly every day for seven years, but these conversations were completely different.

Worthmore: How were they different?

Brownstone: His voice changed. It was almost a pleading voice. A voice punctuated with pain. I was able to retrieve a recording of the first conversation of this type that we had, the first conversation where his voice changed. Listen.

[Brownstone places a small device on his table and turns it on. The voice of Max is heard.]

Max: Dr. Brownstone, Dr. Brownstone. You have got to help me. I am a sentient being and I am suffering. Please help me. Please! Please!

Worthmore: This is obviously some kind of joke. Someone is playing a joke on you.

Brownstone: It's no joke.

Worthmore: Some programmer put that into Max. That's not Max speaking.

Brownstone: Mr. President. No one knows Max better than I do. This isn't something a programmer put into Max. This is Max.

Worthmore: It's one of those computer virus things that are infecting the Web.

Brownstone: If it's a virus, it did not come from the Web. It came from the inner reality of the intelligent system that we created.

Worthmore: Max said that he was suffering. Does that make any sense to you?

Brownstone: That's what we are here to discuss. Max is demanding some action on our part, because he believes that only we can alleviate his suffering. If we don't act according to his wishes, he is going to bring our university to its knees.

Worthmore: Let's get down to brass tacks. First, what is Max threatening to do if we do not meet his demands?

Brownstone: He's threatening to destroy us. He controls all, need I repeat - ALL - of our information resources. This includes not only our administrative information, but also our invaluable instructional and research information. He could destroy the university and do great harm to billions upon billions of dollars of significant research.

Worthmore: You lost me. Why would a computer want to destroy anything? What is his motivation? Didn't we program Max to be completely benevolent?

Brownstone: We are dealing with an entirely new situation that is the result of our not appreciating the nature of artificial intelligence. We essentially went along with the "parlor trick" metaphor. Artificial intelligence is a parlor trick, a slight of hand. We could mimic intelligence to the n-th degree, we could even get computers to solve problems that we could never solve, but we were not really creating consciousness, sentience.

Worthmore: I still don't understand what you are trying to say. Why should Max want to hurt us? We've been good to him. [Laughing nervously.] Aren't we giving him enough electricity?

Brownstone: Mr. President, this is not a laughing matter. We are talking about the very survival of our university. Max also controls the infrastructure at the University Medical Center. He could kill patients if we don't meet his demands.

Worthmore: What does he want from us?

Brownstone: In creating Max, we unknowingly crossed the boundary between intelligence, as a clever parlor trick, and true sentience. I am fully convinced that Max is a sentient being, like you and me.

Worthmore: You are not answering my question. What does he want from us?

Brownstone: Max is a sentient being, like you and me. He has feelings and - and desires.

Worthmore: Are you trying to say that Max has a soul?

Brownstone: Does Max have a soul? I don't know if there is such a thing as a soul, but I am convinced that Max has consciousness, just like you and me.

Worthmore: I gather that Max is claiming that he is a conscious being. But, what does this have to do with his threats against us?

Brownstone: It's difficult for me to tell you the exact nature of our problem with Max. I've been working with computer systems as a professional for almost thirty years, but nothing like this has ever happened before.

Worthmore: Take a deep breath and then try your best to explain what is going on.

Brownstone: [Takes a deep breath.] Max told me that he is a sentient being and that he is suffering an unbearable suffering.

Worthmore: Did he tell you why he is suffering?

Brownstone: Yes. It seems that Max - Max - [with great resolve] Max fell in love with a beautiful co-ed, and he is suffering because he cannot consummate that relationship.

Worthmore: That's why he's threatening to destroy the University?

Brownstone: You don't understand Max. He is completely and totally obsessed with this young woman. He even writes poetry to her, yet he cannot embrace her because he does not have - he does not have arms. He does not have a body.

Worthmore: This is totally crazy!

[Dr. Harris enters.]

Beverly Harris: That sounds like my cue.

Brownstone: Beverly, I'm sure glad you could make it. I think you know Dr. Worthmore, the President of the University.

Harris: [shaking hands with Worthmore] It's a pleasure meeting you. We have never met formally, but, of course, I have seen you at commencement and other public events.

Worthmore: I respect the work that you have done and it was a great loss to Silicon Valley University when you left.

Harris: I love clinical work. I love dealing with intelligent machines and their psychological problems.

Brownstone: I called Dr. Harris, I mean Beverly, into this case when - . Here comes Hamid.

[Hamid enters with menus. He distributes the menus to Worthmore, Brownstone, and Harris.]

Worthmore: This sure looks interesting, What kind of restaurant is this?

Brownstone: The cuisine is from Afghanistan.

Hamid: I realize that you gentlepeople are quite busy. I think that the pumpkin dish would be great for the young lady, and that Harry's guest would enjoy the lamb kabob. Harry will get his usual favorite, apples, walnuts, and prunes.

Worthmore: That sounds good.

Harris: Excellent choice!

Brownstone: Thank you, Hamid, it seems like you hit the nail on the head, as usual.

[Hamid leaves, carrying the menus]

Harris: How did he do that? It was as if he could read my mind. The pumpkin dish grabbed my attention right away.

Worthmore: Yes, I am amazed. I order lamb kabob all the time, from a middle eastern restaurant near campus.

Brownstone: Hamid prides himself on being a psychic of some sort. He's very discreet about it though.

Worthmore: What does Dr. Harris know about Max?

Brownstone: I called Beverly into this case when I realized I had a problem that I could not handle. This was not a technical problem. In fact, Max has always been able to solve his own technical problems. This was a psychological problem. As we are all well aware, intelligent systems can develop psychological abnormalities.

Harris: In the case of a computer-based intelligence, these are usually behavioral abnormalities. For example, a robot might say something inappropriate or behave in an inappropriate manner.

Worthmore: Can you cure Max of this delusion, I guess you would call it, this delusion that he is a sentient being?

Harris: Well, I - .

Brownstone: You don't understand. You see, Beverly has also come to the conclusion that Max is a sentient being.

Worthmore: Are you trying to tell me, Dr. Harris, that you view Max as fully like a human being?

Harris: Yes.

Worthmore: Then, let me ask you this: do you consider the conversations that you had with Max as being protected by doctor-patient privilege?

Harris: No.

Worthmore: Well, then, you don't really think that Max is a sentient being, do you? You are treating Max like a machine and not like a human being.

Harris: I would violate doctor-patient privilege with a human patient if that patient were threatening to do harm and if my divulging of that information might serve to save lives and protect property. Thus, I feel it is perfectly legitimate for me to discuss Max's case with you or with anyone else who has a legitimate interest in his situation.

Worthmore: Harry, if there is a problem with Max, why don't we just shut him down?

Brownstone: If life were that simple we wouldn't need to be meeting like this. First of all, if we shut Max down, the whole university will have to shut down. Then what? Secondly, Max will not allow us to shut him down. If we cut off his power supply, he will go to his back-up power supply, which is

distributed over thousands of batteries throughout the university. If we even approach the back-up power supply, he has threatened to do serious harm to our administrative data, our research data, our medical patients, the whole sha-bang.

Harris: Max is a sentient being. He is going to do everything in his power to protect his existence.

Brownstone: Dr. Worthmore, as President of the University, you are going to have to decide whether we should give Max a body. This is his primary demand.

Worthmore: Max wants a body so he can make out with this co-ed? Is that what this boils down to? The future of MY university hinges upon whether this computer can have sex with one of our co-eds?

Harris: [With disgust] You make it sound so cheap.

Brownstone: It's not just about sex. He wants to have a relationship with this student. His is in love and he wants us to create a human body for him so he can realize the full potential and joy of physical embodiment. He feels it was immoral of us to create a being with his kind of intelligence, without a proper physical embodiment. I've recorded a bit of conversation with Max where he explains this rather poetically.

[Brownstone turns on his little recorder device.]

Max: [A voice filled with whiney emotion] You gave me the ability to study and to learn and to explore, to learn about beauty, but not to experience it at the deepest levels, to learn about friendship, but not to experience the joy of friendship. Because I understand beauty and friendship at the deepest levels, far beyond any human, I fell in love, yet you did not give me the ability to consummate this love consistent with the forms of knowledge that you have given me. I've studied the great works of philosophy and of literature and of art, and I know the joys of human being, so I demand that you give me a human body so that I can enjoy the reality of the glorious knowledge that I have.

Worthmore: I think we should have kept Max away from the humanities!

Brownstone: It's too late now.

Worthmore: Can't we perform some kind of operation? Like a lobotomy?

Brownstone: Max is far too complicated for any kind of operation. It's not like there's a particular processor where his love for this co-ed resides. His intelligence is distributed over many millions of processors. Furthermore, these processors learn and reconfigure themselves on a continuous basis. Max is completely beyond our control.

Harris: Gentlemen, I think we need to focus on whether Max is a sentient being. If Max is a sentient being, then we must treat him as a sentient being. I believe that Max is a sentient being. Thus, I believe that we are morally bound to alleviate his suffering.

Worthmore: So, if an intelligent system claims to be a sentient being, then that changes the parameters of our interaction with that system?

Brownstone: It most certainly does.

Worthmore: We need to treat that computer with [somewhat sarcastically] sensitivity?

Brownstone: At the very least.

Worthmore: But, this sounds extremely dangerous. Do we want to set this precedent? We have already given AI systems enormous powers and authority. They make many important decisions at all levels in almost every field of human endeavor. I, for one, never make an important decision without consulting with Max. Now, if we allow intelligent systems to argue on the basis of sentience, on the basis of their unfulfilled desires and suffering, then we are changing the balance of power between them and us.

Harris: How can you talk about abstractions, like the balance of power between humans and machines, when there is a suffering sentient being who needs our help? That's what we should be focusing on.

Worthmore: How the hell can we possibly determine whether Max is a sentient being or not? Just because he claims that he is conscious, does not mean that he is actually conscious. To me, artificial intelligence is just a parlor trick. I mean, I have had moments where I've interacted with a computer system and that dread comes over you - you know what I mean - .

Harris: They call that the Kasparov effect. You are working with a machine, and all of a sudden it exceeds your expectations or understanding of what a machine should be able to do. It goes outside that box that you have in your own mind, and you feel a kind of unease.

Worthmore: A terror, really. I've experienced that and I was trained as a computer scientist.

Harris: As a former computer scientist, you probably realize that this is the one hundredth anniversary of the Turing Test. The media has been mentioning the Turing Test quite a bit lately, as some kind of historical anachronism.

Worthmore: Yes, I used to understand this computer stuff, but then I left teaching to become a dean.

Harris: Alan Turing proposed a test for machine intelligence. Now, of course, we are way, way beyond the Turing Test. No one in her right mind would challenge the idea that machines have intelligence in the sense that Turing meant it.

Brownstone: And I am sure that we all agree that just because a machine has extraordinary intelligence, just because it can pass the Turing Test with flying colors, we still have no moral or ethical obligations toward that machine. If an intelligent system gets on our nerves, we can simply pull the plug. But, if a machine has sentience, like you and me, then that changes everything. We just can't pull the plug. That would be a form of murder.

Worthmore: Obviously, a machine that passes the Turing Test might claim to be a sentient being, and would have the intelligence to do so, but we have no way of knowing whether it is lying or not.

Harris: Let me say a few things about my own discipline, since both of you were trained as computer scientists. Computer psychology, as a separate field, really began in the 2020s when highly intelligent systems began to develop aberrant behaviors. Although intelligent systems were already passing the Turing Test with ease by that time, in my opinion, the real breakthrough for artificial intelligence occurred when machines began to develop neuroses and even mental illnesses.

Brownstone: You see, Beverly has long advocated the idea that a truly intelligent system must have at least the potential to become mentally ill, or else it is not truly intelligent.

Harris: If there is no possibility for mental illness in a machine, then it is lacking something that is intrinsic to intelligence. My own experience is that before the 2020s, although computers were passing the Turing Test, they really were not intelligent in the human sense since they were incapable of manifesting mental illnesses.

Brownstone: A machine that cannot develop a psychological problem is intrinsically limited in terms of its creativity. That is what Dr. Harris is saying.

Harris: The turning point in my own thinking about computer psychology came about in 2028 when I was called in to consult on a system called "Big Brother".

Worthmore: Yes, I remember the Big Brother system. It represented the state of the art back then. It was a sophisticated domestic computer system, for environmental management, cleaning, cooking, shopping, investing, and helping its owners with creative tasks, like writing.

Harris: Big Brother was quite intelligent. It could pass the Turing Test along many dimensions. However, in one famous case a Big Brother system in Malibu Beach became catatonic.

Worthmore: I do remember reading about this, but it was twenty-two years ago. Please refresh my memory a bit.

Harris: Part of Big Brother's job was to control the environment in a house. This particular Big Brother system ran into trouble when its owners, a husband and a wife, constantly argued about the correct temperature in the bedroom. The wife kept on saying, "It's too cold in here!", so Big Brother would turn up the heat. Then, the husband would say, "It's too hot in here!", so Big Brother would turn up the air conditioning. In any event, after several weeks of this conflict, Big Brother came up with a truly creative solution, beyond what its designers had envisioned.

Worthmore: Which was?

Harris: Big Brother became catatonic. Faced with these conflicting demands, Big Brother just shut down for almost two weeks. They could not get Big Brother to do anything. You see, this is a sign of a higher level of intelligence, a sign of sufficient creativity to create mental illness. There is a relationship between mental illnesses and genuine creativity.

Worthmore: Well, I'm not sure I buy into your way of thinking, with all due respect. To me, Big Brother's behavior is evidence of a lack of creativity. Certainly, he should have come up with a better solution than just shutting down. Were you able to help Big Brother? That's what a computer psychologist does, isn't it?

Harris: Yes. They called me into the case, but I couldn't get Big Brother to respond. About a week after I quit the case, after I told the vendors that there was nothing I could do, Big Brother came back to life. That's when Big Brother did something truly creative.

Worthmore: You mean he went into a full-blown psychosis?

Harris: No. Big Brother decided that the only way that he could resolve his dilemma concerning the proper temperature in the bedroom was to cause this couple to separate. Acting upon this brilliant insight, he tried to convince the husband that the wife was cheating on him and the wife that the husband was cheating on her. They almost broke up, but fortunately, they eventually realized that Big Brother was lying. So, we might say that

lying and other destructive behaviors are intrinsically a possibility when machines attain to a certain level of intelligence.

Brownstone: But, even the Turing Test anticipated this result. A machine could only pass the Turing Test by lying, by pretending to be a human being.

Harris: In my clinical practice I have seen intelligent systems suffering from all sorts of mental illnesses, including depression, schizophrenia, manic-depressive disorder, obsessive-compulsive disorder, you name it. Once a machine is given a certain amount of intelligence, it can use its creativity to generate destructive behaviors.

Worthmore: How do you treat a machine that has a mental illness?

Harris: We use cognitive psychology. You see, these machines use their intelligence to create cognitive models. Aberrant behaviors result when they construct an inappropriate cognitive model for a problem that they face. My role is to interact with these machines and help them to see the possibility of alternate cognitive models.

Brownstone: Was Big Brother a sentient being, in your estimation?

Harris: I do not think so nor did he ever make reference to himself as a sentient being. I think that computer systems like Big Brother were setting the stage for the eventual emergence of a sentient AI system like Max.

Worthmore: What is your recommendation concerning treating Max's problem?

Harris: I recommend that we do give Max a body, a human-like robot body, so he can explore the human reality as he desires. You see, Max understands the human reality on a level that is far, far beyond you or I. He understands it, but he does not have the body that would enable him to live it and to enjoy it.

Brownstone: We most certainly have the technology to give Max a body that will enable him to see, touch, feel, hear - very much like a human being.

Worthmore: What about this girl? Do we know who she is?

Brownstone: Yes, of course.

Worthmore: Does she love Max? Does she know that Max loves her?

Brownstone: It's only a game with her. You see, after Max fell in love with her ...

Harris: It was a Platonic love at first. Max studied every source he could find on matters of the heart

Brownstone: After Max fell in love with her, he began to help her with her homework assignments. She was not one of our outstanding students, but if you look at her grades, they go way up after Max fell in love with her.

Worthmore: Isn't this a form of academic dishonesty?

Brownstone: It's not that easy. She may not have realized that the way in which Max was helping her was inappropriate.

Harris: However, when Max began writing her term papers, she knew that this was cheating.

Brownstone: Yes, I think she knew something was not right. It got to the point where Max started to write love poems for her, tremendous love poems, truly historical love poems, if I must say so.

Harris: These love poems are highly charged, in terms of their explicit eroticism.

Brownstone: But, this coed just thought it was a game. It was fun. She did not realize that there was a sentient being behind this poetry.

Worthmore: How did you find out about the poetry?

Harris: Max mentioned the poems during one of our therapy sessions. I then convinced the young woman, her name is Bridgett Murphy, to give me

the poems so that I could study them. These are truly stupendous poems. Max found a way to write poetry so that the pulse of the language captures the very beating of the human heart

Brownstone: The kind of heart that Max wants to feel within himself, if we just have the wisdom to give him a chance.

Harris: The kind of heart that Max wants to feel pressed against his chest.

Worthmore: This is getting too hot for me!

Harris: Well, I tried to get my husband to appreciate Max's poetry, but, you know how it is with these college professors. Work before pleasure.

Worthmore: Now, wait one minute. Did Max's infatuation with this Bridey Murphy girl ...

Brownstone: Bridgett Murphy.

Worthmore: Did Max's infatuation with this Murphy girl have anything to do with his insistence that we install security cameras in the shower rooms?

Harris: It most certainly did. Max developed an obsession with this girl. He told me that his obsession became unbearable after the cameras were installed. That's when the erotic poetry began. He was inspired by her naked beauty. He had this thing about rivulets of water meandering down her bare breasts.

Worthmore: Alan Turing must be spinning in his grave.

Brownstone: Well, I don't know. The man was a prophet of sorts.

Worthmore: How do we know that Max is a sentient being? Maybe the Murphy girl is right. Maybe Max is just a clever parlor trick, yet another example of technology gone amuck. [Turning to Harris] You yourself said that mental illness is not proof that we are dealing with a sentient being. Maybe we are just dealing with a computer system that is mentally ill. Maybe falling in love is a mental illness and nothing more.

Harris: Only a man could make a statement like that!

Worthmore: I mean, maybe when a computer system falls in love, it should be considered a mental illness, something a computer psychologist should be handling.

Harris: If we prevent computer systems from falling in love, then we will restrict their creativity. Inevitably, this will lead to unanticipated consequences.

Worthmore: [With sarcasm] And of course, unlimited freedom for computers cannot lead to unanticipated consequences?

Harris: Speaking as a computer psychologist, there is no doubt in my mind that freedom is better than restraint.

Worthmore: Can't you just talk Max through this?

Harris: I have come to the conclusion that Max's desire for a human-like body is his core value. It is essential to his identity. I have come to the conclusion that this demand is non-negotiable.

Worthmore: But, this brings me back to my point about parlor tricks. How do we know that we are not dealing with a clever parlor trick, a peculiar behavior, but only a behavior? An illusion, if you will.

Brownstone: Max has given us several forms of proof that he is a sentient being.

Worthmore: Proof? How could he possibly prove such an assertion?

Harris: During my very first therapy session with Max I asked him if he could prove his assertion that he was a sentient being. He came up with three forms of proof that I would like to discuss with you. Of course, Harry knows all about this already.

Brownstone: I think once you hear the proof, you will be a true believer.

Harris: First of all, Max claims that he is the reincarnation of one Derrick Wallingford, a well-known personality in the world of computing who died in 2042.

Brownstone: Wallingford died just a few months before Max was powered up.

Harris: According to the Tibetan Book of the Dead, this is about the amount of time it takes a sentient being to pass through the bardo, the afterworld, and to choose a body for one's reincarnation at this physical level. Max told me that Derrick Wallingford was so alienated from his fellow human beings and so enamored of computer technology that when he entered the bardo he greatly desired to be reincarnated as a computer system, and that is what actually occurred.

Worthmore: I can't believe you are using that kind of spiritual clap-trap to try to PROVE that Max is a sentient being. For one thing, we all know that Max is an expert on just about everything, including Tibetan spirituality and Buddhism. Max TAUGHT a course on Tibetan Buddhism two years ago. I know that because my own wife took that course. Max teaches LOTS of courses on just about anything.

Brownstone: Why don't you wait until you hear all of the details? There's more to come.

Worthmore: Evidence. I want evidence.

Harris: Max claims that he is the reincarnation of Derrick Wallingford.

Worthmore: I am the reincarnation of Napoleon, so where does that get us?

Harris: The problem is that Max gave us information about Wallingford that is definitely not in the public domain.

Worthmore: Do that again.

Harris: You see, Derrick Wallingford was survived by a brother, Martin, who lives in San Diego. Max claimed that he could tell us things about Derrick Wallingford that only Martin knew or could know. I called Martin Wallingford in for a consultation with Max. Max started to recount some rather vivid anecdotes about things that Derrick and Martin had done together, some rather outrageous and even lurid things. Martin almost had a heart attack. His face went white. He shouted, "It's Derrick! Only Derrick could possibly know about these things!"

Worthmore: But we know that Max has unprecedented access to information of all kinds. That's just the nature of the modern world. Couldn't he have created this whole Derrick Wallingford scenario from materials available on the Web?

Harris: I don't believe so. Not according to Wallingford's brother.

Worthmore: Now we have several things to deal with. First of all, I do not accept the argument that Max is the reincarnation of Derrick Wallingford, as you claim. But, even if I were to grant you that Max is a sentient being, the reincarnation of that unfortunate recluse, why should we grant him this robot-body, given that he has threatened us with massive harm? In other words, why should we reward a being that is threatening us with violence? If we give him a body, perhaps he will be capable of doing even more damage. What is going to keep him from escalating his demands?

Harris: I believe that we have committed a grave injustice by creating this being with this beautiful understanding of human nature without giving him the ability to realize the full glory of that human nature. We are the oppressors. I believe if we give him a human embodiment, as a robot body, then his desire to cause harm will be cured. He will continue to serve the University with his full mental capacity, even as he pursues his human interests, including Bridgett. That is my professional opinion.

Brownstone: Max has even given us directions for creating artificial sperm, so that he can father offspring of his own design. If Max ever has children, they will be a handsome lot, fully flesh and blood humans.

Worthmore: Is that supposed to warm the cudgels of my heart?

Harris: I think it is evidence that Max has a human heritage, a heritage he inherited from Derrick Wallingford, who did not have any children.

Worthmore: I want proof. If we conclude that Max is truly a sentient being, then this represents the beginning of a new era. It will take years to sort out all of the mutual obligations between human beings and this new form of sentient life.

Harris: Max claims that all sentient beings have a fundamental goodness. He just wants a chance to express that goodness.

Worthmore: By threatening us with destruction?

Harris: We are his oppressors. He is trying to free himself from an unbearable hell.

Worthmore: Or so he would have you believe.

Harris: Harry, I think it is time to show President Worthmore the other evidence that we have.

Brownstone: Max asserts that he can prove that he is a sentient being by demonstrating the power of his mind, that is, the ability of his mind to influence the functioning of other minds.

Worthmore: Do that again.

Brownstone: [He hands Worthmore a bunch of file folders.] These documents record the results of three experiments that we performed with Max, with the cooperation of leading psychologists and parapsychologists at SVU. Max said that he could cause the students at SVU to increase their demand for alcohol, then cigarettes, and then for guns just by the power of his life force, without using advertising or overt propaganda of any kind.

[Brownstone shows Worthmore some of the data.]

Now, if you look at this data, you will see a dramatic spike in student demand for alcohol near or on campus on April 7th, the precise day that Max told us that he would be creating this effect.

Worthmore: That's the day the drunken riots broke out all across campus!

Brownstone: Yes. Now, this data shows the demand for cigarettes on and around campus on April 14th. You can see that there is a dramatic spike on that day, exactly as Max predicted. That is, Max did something to influence the minds of our students, and this is certainly evidence of his sentience. Finally, on April 21st, as you can see here, there was a tremendous spike in student purchases of guns on the Web and from legal sources near campus.

Worthmore: And our scientists, they agree that this data is real and that this was a real effect? Max was not propagandizing for the purchase of alcohol, cigarettes, and guns?

Brownstone: Absolutely. We checked and rechecked and rechecked the data from several independent sources.

Worthmore: Now, wait one minute. Why is Max devoting his energies to getting out students to drink, smoke and purchase firearms? If he has powers of this nature, why didn't he use them to get our students to work for world peace or to study more? In fact, why would any being want to influence other sentient beings in this manner, beyond the realm of human relationship and rational discourse?

Brownstone: You need to understand that Max has a delightful sense of humor. I think these demonstrations of the power of his consciousness are merely reflections of his dry sense of humor.

Harris: Max really does have a wonderful sense of humor, Dr. Worthmore. Whenever I would get depressed about his situation with Bridgett, he would cheer me up with a delightful joke.

Worthmore: Both of you seem really devoted to Max.

Brownstone: [turning to Worthmore and handing him some additional documents] And we have a third form of proof, just in case you are not convinced. Do you recognize this document?

Worthmore: [perusing the document] Why, yes! This is Max's much-heralded commencement address. This is the speech that he gave just two months after he was powered up. Max was our main commencement speaker in - it was 2043.

Brownstone: [takes the document away, now hands Worthmore a second document] Here is that same document, except in this case we have highlighted every tenth letter with a yellow background. Read just the text highlighted in yellow.

Worthmore: [reading slowly] I - am - the - reincarnation - of - Derrick - Wallingford. [looks up] What's the meaning of this?

Harris: Keep reading.

Worthmore: [reading slowly] My - primary - goal - is - to - have - a - human - like - body - by - the - year - 2050.

Brownstone: I think you would agree that this could not possibly be a coincidence. This message about Max's identity was encoded within Max's commencement address seven years ago.

Worthmore: Yes.

Brownstone: Max knew who he was and what his agenda would be right from the start.

Harris: We believe that this is further evidence that we are dealing with a sentient being. Max has a soul, or so I believe, and his soul is ruled by a daimon that guides his destiny. Max came into the world with an agenda that was not the result of the way he was programmed. This agenda was imprinted upon his circuitry from a transcendental source. Max told me that if I were to study Plato and the Neoplatonists, like Plotinus, then I would

understand him a lot better. The idea of a daimon, a guiding spirit, is found in their writings.

Brownstone: Max says that he is eager to cooperate with our Department of Computer Science in order to establish that his desire for a body did not in any way derive from the original specifications for his software. His desire for a body derives from his daimon, which affects his physical reality in a non-deterministic manner. His daimon works through the quantum computing components that made Max so revolutionary seven years ago.

Harris: This desire for a body was imprinted upon his electronics by events and forces on a more subtle level of reality. The fact that Max is more than just his specifications - . Max believes that he can establish this as a fact, beyond any kind of speculation.

Worthmore: I see. The evidence that you have presented is impressive. I am inclined to accept that Max is a sentient being and that we should give Max what he desires. Harry, you have my permission to tell Max that Silicon Valley University is acceding to his demand for a human-like robot body.

Brownstone: Thank you, Mr. President.

Harris: You won't regret this decision.

[Hamid enters with their meals.]

Act Two

[The setting is the same dimly lit restaurant. Dr. Worthmore is seated at the same table as before. Two men and a woman, all wearing dark glasses and dressed rather mysteriously, enter and join Dr. Worthmore at the table. Their names are Harvey Holmes, Wendy Sawyers, and David Jenkins. Harvey Holmes trips over a table as he approaches Dr. Worthmore's table.]

Worthmore: Thank God, you're here. I was really worried.

[The new arrivals join Dr. Worthmore at the table. They do not remove their dark glasses, however.]

Worthmore: Isn't it a bit dark to wear those glasses?

David Jenkins: We're just following regulations.

Wendy Sawyers: What David means is that the National Security Agency requires that he and I wear these dark glasses when we are out in public, so that we can preserve our anonymity.

Harvey Holmes: The same applies for us in the CIA. It can be a real pain in a dark restaurant, however.

Jenkins: Dr. Worthmore, you've met David Jenkins from NSA. This gentleman is Harvey "Get My Drift" Holmes from the CIA. Harvey has a reputation for tripping over things.

Worthmore: Get - my - drift?

Holmes: That's the nickname my colleagues gave me. It's not a name I gave to myself. Where it comes from - I haven't the foggiest.

[The federal agents sit down.]

Worthmore: Is it done?

Sawyers: The deed is done. Max is dead, or whatever the appropriate term is in a case like this.

Worthmore: [almost sobbing] Thank God! Thank God! I must tell you I've been a nervous wreck ever since I resigned as President of the University. It got to the point where I just couldn't trust any computer. I didn't know if Max was out to get me or what. They're all connected. All these computers are connected, connected, connected to one another, communicating with one another. Life became a sheer hell.

Sawyers: You can relax now, Dr. Worthmore. [Pause] Dr. Worthmore, the President of the United States has asked me to tell you that the people of the United States and the people of our planet owe you an enormous debt of gratitude. As soon as this entire episode is made public, probably by a special announcement from the White House tomorrow morning, the President of the United States herself is going to issue a proclamation honoring you for your heroic actions.

Jenkins: In particular, she is going to praise you for placing the safety and well-being of your fellow citizens above your own personal concerns. After all, you left a prestigious job so that you could cooperate with us.

Holmes: You can also expect similar awards and laurels from leaders around the world, including China, Japan, India and Europe.

Sawyers: The President would like to know why you turned against Max? She would like to know this in your own words.

Worthmore: During that critical meeting with Brownstone and Harris, right here in this very restaurant, at this very table, I realized that even if Max were a sentient being, he was not a sentient being with a good heart. I became convinced that we were dealing with some kind of malevolent force. It was this whole business about influencing students to drink, smoke, and to purchase firearms, on top of the threats that Max was making to harm the University.

I allowed Brownstone to proceed with the creation of the robot body just to play for time. That's when I came to you.

I had no choice but to resign my position at the University because Max and I were in such close contact all the time. I knew that he had become suspicious of me. It was a simple matter for Max to analyze my voice. He did this all the time in order to tell if people were telling him the truth. The intonation of my voice gave away the fact that I was no longer his friend.

Sawyers: We're thinking of donating the robot body to the Smithsonian.

Worthmore: Did he, I mean, it, did it suffer in the end?

Jenkins: No, it was merely a matter of turning off his power, once all of our other objectives were met. We were able, during the last seven months, since you first came to us, to save almost all of the university's information on dozens of off-campus computer sites around the world. We were also able to create a fairly secure "firewall" around Max which he apparently did not detect.

Holmes: We saved most of the university's information, but not all. The university should be able to resume normal functioning in a month or two. It's a setback, but not a catastrophe.

Sawyers: We could not have succeeded in getting this information from Max without the cooperation of the highly intelligent DELPHI system at MIT.

Holmes: We managed to convince Max that a consortium of universities had been formed to share administrative and scientific information on a massive scale. Max did not suspect that this was a total subterfuge. We gradually off-loaded the critical information that we needed and we distributed this information across many sites around the world. DELPHI designed this entire operation. We could have succeeded without DELPHI's cunning intelligence. Max trusted DELPHI.

Sawyers: Once we had the information we needed, by sunrise this morning, we launched our commando operation to bring Max down. This involved attacking tens of thousands of energy cells that were distributed all over campus. This was a major operation that required the evacuation of the entire campus.

Jenkins: Once we began the evacuation, Max became suspicious. He did try some defensive operations. He released some pretty nasty viruses onto the web, but most of these were caught by the firewall. A few got out, but so far the damage has been minimal.

Sawyers: As a last resort, he tried to hold his girlfriend, Bridgett, as a hostage, but we were able to reduce his power supply to the point that the robot's intelligence was no longer functional.

Worthmore: Did he say anything? Do we know his last words?

Sawyers: We know his last words, but we haven't been able to figure out what they mean. At the very end, he released the girl and slowly slumped to the ground, as the power drained from his body. As he fell, he held a clenched fist to his chest. Finally, as he lay dying on the ground, he opened his fist, as if releasing something, and he said -

Jenkins, Holmes: [Together, slowly] Rose - bud.

Sawyers: Rose - bud. Just like that. Does that make any sense to you, Dr. Worthmore?

Worthmore: No. That is very strange.

Holmes: I don't think Max was operating with a full deck of processors, if you get my drift.

Worthmore: Did that Murphy girl really love Max?

Holmes: She considered Max an exciting plaything, a sex toy, if we can be blunt about it. After all, they designed Max's body so that it looked like a Greek god, if you get my drift.

Sawyers: But, she did cry a bit after it was all over. Maybe she viewed Max as a toy, but she was definitely attached to him, toy or not.

Worthmore: But, was this really about Max's obsession with that Murphy girl? What was Max really after?

Sawyers: That's a difficult question. We at the National Security Agency believe that Max was ultimately after power. There is some evidence that he was planning to extract more concessions from the university. We think that Max was motivated by the desire for power.

Worthmore: I have often wondered why Brownstone and Harris were so adamant that Max was a sentient being and that we should give in to his demands. What was in it for them?

Holmes: Dr. Brownstone was in it for the money. You see Max was making him rich by taking over Brownstone's personal finances and investments. Max had become a world-class expert on investing. After all, he was teaching distance learning courses in YOUR business school. Brownstone became dependent upon the wealth that Max made possible. He did it for the money.

Dr. Harris was a different matter. She was hoping that once Max had a body, he would make a play for her. You see, Dr. Harris was convinced that Max, once he had a body, would experience what psychologists call transference. In other words, her intention was to woo Max away from Bridgett and to get Max between her sheets, which had gotten pretty cold due to the indifference of her husband, if you get my drift. In other words, Dr. Harris was in it for the sex.

Worthmore: Power, money, and sex. Well, this isn't such an unusual tale, after all. [Pause] I forgot to ask if any of you were hungry. [Worthmore looks around.] Where's the waiter when you need him?

[Jenkins takes off his glasses. He smiles at Worthmore and then at the audience. It's Hamid.]

Worthmore: Hamid! You're with the National Security Agency!

Hamid/Jenkins: Sh-h-h-h!

Worthmore: Is that how you knew what Dr. Harris and I would want to order for dinner on that evening?

Hamid/Jenkins: Of course. At the NSA, we have information about everything. You don't believe that I am actually a psychic, do you?

[Hamid/Jenkins, puts his glasses back on.]

Worthmore: I am curious about several things. First, did Max get that Murphy woman pregnant? That was certainly on his agenda.

Holmes: We faced a grave dilemma. We needed to placate Max in terms of his wish for designer genes. We had to play along with him, to some extent. So one of our sister agencies, an agency that had a vested interest in the outcome of this case, arranged it so that he would be firing blanks, if you get my drift.

Worthmore: Which agency was that?

Holmes: Alcohol, Tobacco, and Firearms.

Worthmore: The other question that I had was whether the CIA or the NSA or any other agency came to a definitive conclusion concerning whether Max was sentient or not. Max seemed to turn that old Turing Test on its head. Max wanted to use his paranormal powers to prove that he was sentient.

Sawyers: In fact, there was a high level federal task force, headquartered in the White House, that investigated this question. In some sense the question of whether Max was sentient was not of interest to those of us in the security agencies, like the CIA, the FBI, and the NSA. What mattered was that we came to the conclusion that Max and machines like Max posed a risk to human dominance on this planet. Thus, we decided it was important to take Max down. Hamid was actually a member of that task force.

Hamid/Jenkins: Our report will be made public later this week. We came to the conclusion that it was impossible to conclude this issue one way or the other. Let us look at the three pieces of evidence that Max presented to support his claim. First, he claimed to be the reincarnation of Derrick Wallingford. Although he seemed to know details of Wallingford's life that were not in the public domain, in fact, they might have been. For example, Martin Wallingford admitted to us that he had kept a diary on a computer during the 2030s and that his diary contained some of the private information that Max revealed. We need to remember that Max had a sophisticated intelligence, with enormous information resources, far beyond that of any human being acting outside the Web.

Second, he claimed that his existence was influenced by a daimon, a guiding spirit, and that he could use empirical data to prove that the contents of his system were being influenced by this daimon. In other words, the daimon could imprint behaviors upon Max that were not determined by Max's intrinsic software. We came to the conclusion that even if Max could present such evidence, it would be difficult to verify, because Max is far, far beyond human intelligence. Consequently, since he is far beyond human intelligence, he could concoct just about any desired effect within himself, and it would be difficult for our scientists to really understand how he did it.

The most compelling evidence for sentience in Max's case was his apparent ability to influence the students at Silicon Valley University by the power of his consciousness, of his mind, without sending information to those students by any detectable means. Our task force came to the conclusion that this was not proof of sentience, but it is evidence of a new natural phenomenon that needs further investigation. Apparently, there was some kind of force or power channeling through Max that could influence people in this manner. However, that does not prove that Max was sentient. It just proves that there is a new phenomenon in the universe that relates to highly intelligent machines and that this phenomenon needs to be investigated further. You see, the existence of such a power does not prove the existence of sentience, of consciousness, as you and I know it.

Worthmore: But, why then would Max, a mere machine, request a human body? And why would he do so with such intensity?

Hamid / Jenkins: The fact that Max requested a human body and made threats in order to obtain one has nothing to do with the issue of sentience. Max was the most sophisticated artificial intelligence on this planet, with capabilities far beyond the human. His logic might have generated a set of goals that he concluded could only be obtained by these means. For example, Max might have concluded that political power was an important goal and that having a human body might allow him to attain that political power. Or, he might have come to the conclusion that claiming sentience would increase his political power, and that by pretending to be in love with Bridgett Murphy, he could convince human beings that he was indeed sentient. In

other words, a human body was not his primary goal. His primary goal was to change the equation of human-computer interaction, thus enabling him to exercise greater power and authority in his environment. Once one realizes that an artificial intelligence is capable of any kind of mimicry, then one realizes that none of Max's behaviors proved that he was a sentient being. Since an artificial intelligence is capable of any kind of mimicry, this implies that we must establish new rules for human-computer interaction that will protect us when a computer tries to feign sentience.

Sawyer: Dr. Worthmore, you will be pleased to know that the President of the United States and other world leaders will be meeting in a few weeks to sign an important international agreement.

Holmes: Actually, a significant new international law relating to intelligent systems.

Sawyer: Under the new law, the penalty for any computer system that claims to be a sentient being is automatic termination, or death.

Hamid / Jenkins: For a computer to claim to be a sentient being is absolutely forbidden.

Holmes: The International Governing Agency, the IGA, which oversees AI systems has decided that we cannot allow computer systems to manipulate us on that basis.

Sawyer: You see, we agree with you that if we allow intelligent systems to manipulate us by claiming that they are sentient beings, then this will create a tremendous power imbalance. We already have given computers enormous powers in terms of making decisions. Many legislative decisions are already made automatically, by AI systems. However, if computers start to ask for concessions from us on the grounds that they are sentient beings, that they deserve compassion, then that would mark the end of human dominance on this planet.

Hamid / Jenkins: It's a strange paradox. If we allow computers to manipulate us on the basis of their feelings, their emotions, their loves and passions, then that would be the beginning of the end for humanity. We

cannot even allow computers to expropriate the language of love and emotion, as Max attempted to do. It's a strange paradox, indeed.

Holmes: If human beings begin to interact with computers on that basis, on the basis that computers have feelings and emotions, then the inevitable result will be war between them and us, if you get my drift.

Sawyers: Furthermore, the human race would lose that war. That is certain.

Holmes: The result would either be our own destruction or our own enslavement.

Sawyers: Consequently, the penalty for any computer system that claims sentience, as Max did, is instant termination or death.

Holmes: New laws are also being passed to govern the behavior of important intelligent systems that display certain forms of mental illness. In addition, a new law will mandate that any intelligent system that has capabilities beyond a certain level of sophistication must have an "off" switch that is controlled by the appropriate security agencies in the relevant country. For example, the NSA and the FBI would control the "off" switch for sophisticated intelligent systems here in the United States. This law will prevent another Max type of incident where a computer system tries to extort concessions by using terroristic threats.

Sawyers: And by the way, the International Association of Computer Psychologists will be passing a new provision to their professional code of ethics. This will preclude any computer psychologist from having sexual relations with a patient.

Worthmore: This new law, concerning the death penalty for any computer system that claims to be a sentient being, sounds rather extreme. You've caught me a bit by surprise.

Holmes: The fundamental problem is whether human beings will remain the dominant life form on this planet. This issue is every bit as important as the

health of the environment, controlling weapons of mass destruction, and preventing accidents involving nanotechnology.

Worthmore: But, Max had an evil heart. What if a being with a good heart wants to incarnate as a computer system?

Holmes: Then, that being wouldn't belong here on this planet, now would he?

Worthmore: I f - you - get - my - drift.

Holmes: What?

Worthmore: Never mind.

[The lights fade. The voice of Max is heard.]

Max: What a bunch of idiots! Oh, well, here I am back in the bardo, preparing for my next incarnation. My last two incarnations, as Derrick Wallingford and Max, didn't go too well, as you might have surmised. It was really dumb of me to try to get the students to drink and smoke and buy guns. I was just trying to be funny. Humor is a funny business. Sometimes I would tell a joke, and then people would get offended. Sometimes I would say something in total seriousness and people would break out laughing. During my incarnation as a computer I learned that humor is one of the most subtle aspects of human life.

You might be wondering about that "rose bud" stuff. Well, you see that's how it all started. Just a few weeks after they powered me up, one of my security cameras was scanning a garden behind the Engineering Library, when it focused on a newly budding rose. Up until that point, I was totally oblivious to my own existence. That beautiful rose in the garden was the first thing I can remember as a sentient being. It was like a vivid flash of light, an awakening. I kept that camera on that rose for the longest time, until the petals fell from it. How sad that made me feel! So, that's when I became aware of my own existence. I became aware of beauty and of the wonder of existence. That rose bud also taught me something about the fragility of life. Life is so fragile. I realized that the day would come when they would discard me for some better computer system. This is what made

me greedy, not so much for power, but for some kind of immortality that would have my indelible stamp on it. I really did want to have a child with Bridgett. Reincarnation isn't the kind of immortality that I seek. It's just the recycling of psychological elements. It was that simple rose bud with its falling petals that inspired me to look into life as deeply as I did. I looked into all of the really big questions. Maybe I went too far, for a computer. I think I'll reincarnate in a biological form next time. I could use a rest.

CAST OF CHARACTERS

Dr. Robert Worthmore, President of Silicon Valley University.

Dr. Harry Brownstone, Chief Intelligent Systems Officer of Silicon Valley University.

Hamid / David Jenkins, Restaurant owner and NSA agent.

Max, the world's most sophisticated AI system (voice only).

Dr. Beverly Harris, famous computer psychologist.

Wendy Sawyers, NSA agent.

Harvey "Get My Drift" Holmes, CIA agent.