

Penn State Journal of Law & International Affairs

Volume 7
Issue 3 *Symposium Issue*

April 2020

Autonomous Systems and the Meaning of Humanity

David Danks

Noreen Herzfeld

Amy Pritchett

Matthias Scheutz

Follow this and additional works at: <https://elibrary.law.psu.edu/jlia>



Part of the [International and Area Studies Commons](#), [International Law Commons](#), [International Trade Law Commons](#), and the [Law and Politics Commons](#)

ISSN: 2168-7951

Recommended Citation

David Danks, Noreen Herzfeld, Amy Pritchett, and Matthias Scheutz, *Autonomous Systems and the Meaning of Humanity*, 7 PENN. ST. J.L. & INT'L AFF. 195 (2020).

Available at: <https://elibrary.law.psu.edu/jlia/vol7/iss3/6>

The Penn State Journal of Law & International Affairs is a joint publication of Penn State's School of Law and School of International Affairs.

Penn State
Journal of Law & International Affairs

2020

SYMPOSIUM ISSUE

**AUTONOMOUS SYSTEMS AND THE
MEANING OF HUMANITY**

Moderator: Ben Johnson

*Panelists: David Danks, Noreen Herzfeld, Amy Pritchett, and Matthias
Schantz*

Ben Johnson: Alright, so we're going to get back into this. This has been a great day full of broad, wide ranging topics. We thought it would be really useful to bring this to a conclusion with a very narrow question, which is the relationship between autonomous systems and humanity. With that narrow thing, we are going to focus like a laser beam with four excellent panelists. I will introduce them briefly, and then turn them loose to enlighten us all.

First, to my immediate left is David Danks, who's the L.L. Thurstone Professor of Philosophy and Psychology at Carnegie Mellon, where he heads the Philosophy Department as affiliated with the Institute for Strategic Analysis. He works with the intersection of philosophy, cognitive science and machine learning. His earlier work in computational cognitive science led to a book in 2014 of MIT Press, *Unifying the Mind: Cognitive Representations as Graphical Models*. You can find it on Amazon, I checked, it is still there.

David Danks: Really?

Ben Johnson: Indeed.

David Danks: Alright.

Ben Johnson: He is currently working on autonomous vehicles with an eye toward making sure that our practices remain human-centric, rather than technology-centric. After this is over, I am going to be picking his brain about Beijing epistemology. If you have questions about that, he will also take those as we go forward.

Noreen Herzfeld has perhaps the coolest job I have ever heard of. She is the Nicholas and Bernice Reuter Professor of Theology and Computer Science at the college of Saint Benedict and Saint John's University. She has written several books that are also available on Amazon. You should go and check it out at the intersection of technology and religion. Her current work focuses on what AI does to human relationships, both human-to-human relationships, but also whether we can have authentic relationships as humans with AI.

Amy Pritchett has the Department of Aerospace Engineering here at Penn State after we stole her away from Georgia Tech, or as I like to call it, the North Avenue Trade School. She is the past Director of NASA's Aviation Safety Program. She is deeply interested in the intersection of automated technology and human performance, especially in the aviation arena. She is a true visionary when it comes to theorizing and integrating concerns about safety into the training and practice of engineering.

Ben Johnson: Since getting to know Amy I have been a beneficiary of her amazing culinary talents. I

have also grown increasingly weary of ever getting on an airplane. Alright, Matthias Scheutz, did I say it correctly?

Matthias Scheutz: Perfect.

Ben Johnson: Yes. I was very nervous about that one. Professor of Computer Science at Tufts and the Director of the Human-Robot Interaction Lab, where he researches topics related to AI, human-robot interaction and also teaches in the Philosophy Department. He got PhDs in both Philosophy and Cognitive Science/Computer Science. He is a lot smarter than I am. He also has a book on Amazon, but I'm particularly interested in his recent work on what he calls "MacGyver Problems." Not because I understand any of it, but because I am old enough to remember and love the original TV show. With that, I am going to turn it over to David, and I will give you two-minute warnings and then cut.

David Danks: Okay, great. Thanks to CSRE and JLIA for inviting us, and having this really amazing one and a half days. Thanks to all of you, I was going to say for sticking around at the end of a long day, and then I realized it is only two o'clock. This is what happens when you start at 8 AM, is that it's a long day and it is still the middle of the afternoon.

David Danks: I wanted to talk a little bit about this issue of autonomous technologies and the meaning of humanity, this nice, vague, high level issue. Ben did a wonderful job in some emails with the four of us, in which he laid out some questions that he thought might be interesting. One of them really jumped off the page to me, which is the question of, what is our current

relationship with machines doing to us as people? I think this is a question that comes up a lot. We even just started to hear it at the end of the last session, about how children, their cognition is being changed by virtue of the interactions that they're having with their machines. I think we see this also in a lot of the discussions from the morning panels.

David Danks:

Why do we care about having trust in the AI? Why do we care about having explainable AIs? Because of our engagement with the AI. How is our relationship, the nature of our cognition changed by virtue of our relationship with the AI, with the autonomous technology? I think that this is the way that this question often goes when somebody poses it. In the true contrarian spirit of a philosopher, I want to say I think that's actually not the way the question should go. I think instead we should focus on something that is utterly essential to our humanity, namely we are all social beings.

There is ample psychological evidence, whether from health psychology, developmental psychology, social psychology that we humans suffer when we're deprived of social relationships with one another. That a critical part of the development of the next generation is their successful interactions with one another. Anyone who's a parent knows, kids sometimes need to get into arguments with each other. They need to learn how to resolve conflict. These social bonds that bind us together are absolutely critical, not just to the successful function of the society, but also to our very humanity.

David Danks:

It seems to me we need to ask the question about, not just how our relationships with machines are changing us as individuals or us in relationship with the AI, but also asking: how does our relationship with the machines impact our relationships with one another? How are our social interactions being changed by virtue of the ways that we interact with the systems? There is quite a lot of discussion of that right now with regards to social media. It is sort of ubiquitous if you pick up *The Economist* or *The Wall Street Journal* or pick your favorite somewhat high minded public popular press. There are discussions about how social media is the end of love or the end of relationships.

There is actually surprisingly little, though, looking at how engagement with autonomous technologies or AI systems can change our social relationships with one another. Let me try and give you an example of the kind of thing that I have in mind. Suppose you go to your doctor, some people do I hear. You go to your doctor, and there's a series of tests and your doctor looks at you and says, "Great, I'll let you know what's wrong with you when the AI tells me."

Now, one of the core tenets of western healthcare is the bond of trust and shared decision making between patient and doctor. What does that do to your relationship of trust when your doctor becomes an information broker to you? When all the doctor is, is merely a go between, between a diagnostic AI system and you. At that point what do you need the doctor for, you might start to wonder.

David Danks:

Now you might think this is far off in the future, but if you talk to anybody who works with healthcare companies, they will tell you that all healthcare providers and insurance companies are already looking quite seriously at when they will flip the switch and mandate the use of diagnostic AI systems for doctors. It's basically the exact same argument as you get throughout epidemiology, vaccines and so forth. If these things are better at doing the job that humans are doing for diagnosis, then even though there's a misdiagnosis rate, if it's small enough, we need to mandate it otherwise we open ourselves up to medical malpractice. I'm not a lawyer, please don't quiz me on that part, I'm just reporting what I hear.

Now, imagine instead you come into your doctor, and your doctor says, "Let's figure out what matters to you. I have this system here that is going to be helpful in terms of trying to diagnosis what biochemically might be the problem with you. That's going to enable me to focus on what your values are, what matters to you, and what treatment options to use of the fifteen that I've actually had the chance to learn about. I'm not worrying about staring at lab tests, but what might be the best for the situation in which you find yourself." Now the innovation of the diagnostic AI is increasing your trust in your doctor, is increasing and strengthening the social relationship that you have.

David Danks:

Just as we often think about AI systems as augmenting or replacing or impairing individual functions—it's replacing me on the assembly line—we can equally well think about AI systems as augmenting, impairing or

replacing human-to-human social relationships. There are ways that if we do it right, AI systems can substantially improve our relationships with one another. There are ways if we do it wrong, that they can substantially impair those relationships. I think it's imperative that we not just think about human AI connections when we think about how our humanity might be changed, but think about the alterations in this core feature of our humanity, which is the relationships with one another.

David Danks:

I'm not trying to suggest a pessimistic outcome, I'm actually a techno optimist, just a realistic one. That if we don't take some measure of ownership in the development and deployment of these technologies, we're not likely to end up in a world where our human-human relationships are supported rather than undermined.

I just want to close with just one or two minutes about a connection with security, because we have heard a lot about things like autonomous weapons. I talked about healthcare, what do these have to do with one another? I think one of the things that's notable is if you talk to people in the security field, how many of them will emphasize the importance of interpersonal relationships. How much they will talk about crisis averted, because two people knew each other on opposite sides of the fence. It's not the *Journal of Law and International individuals*. It's the *Journal of Law and International Affairs*, affairs in the sense of relationships with one another.

- David Danks: This building houses the School of International Affairs. We think of international security as being intimately tied with international diplomacy, an engagement in politics. I think that that just reveals the importance of thinking about not just the way that these AI systems can destabilize an individual military, whether the adversary's use or our own use, but also about the impacts on security as it emerges from the deeply personal human-to-human interactions and relationships that are really at the core of a lot of our lives. With that, I will hopefully have allowed some extra time that I'm sure my panelist could use more wisely than I. Thanks.
- Ben Johnson: Noreen.
- Noreen Herzfeld: I'm not sure about the wisely part, and since it's late, I decided to show a lot of pictures. That usually works with my students when you have a class that's late in the day. When Admiral Houck said someone that said to him, "Well, I really don't know why I'm here," that was me theologian. This is not the usual setting where I talk. The question, though, that I would like to address is, what do we want from AI and why? Can traditional theology tell us anything about the answer to that?
- Noreen Herzfeld: I want to come back to something that Ben Johnson said this morning, and that is that we learn awful a lot about ourselves, about what we're looking for, what we're aiming at and maybe where we're going by looking at the stories that we tell ourselves. If we look at the stories we tell ourselves about AI, well we've got *Her*, we've got *Ex Machina*. You might

notice that following our previous presenter, they're all about relationship.

Noreen Herzfeld:

Real AI looks more like a Roomba most of the time or like a factory arm. It doesn't look the stories that we're telling. When we think about what it is that we're actually looking for, I go back to Genesis one, to "in the beginning." Here is the verse from the Bible that says we were created in God's image and I'm particularly interested in the parts that I bolded, "Let us make humankind in our image, according to our likeness, . . . to have dominion . . . And God created humankind in His image. In the image of God he created them." When theologians look at this verse and they say, "Okay, we're in the image of God, but what does that mean?" The verse doesn't spell it out.

When we think about AI, we think, "God made us in God's image, we are trying to make AI in our own image." What I think we are actually doing is standing in the middle and projecting in two directions. We're looking at what we might share with God, then what we would like to share with the computer. When we look in those two directions, we see three possibilities. These are the three ways theologians have typically understood us as being in the image of God – as reason, function and relationship. If you think about those three, they fit AI very well indeed.

Noreen Herzfeld:

We start with reason. This was the first way theologians such as Augustine and Aquinas, even Calvin and Luther, thought that we imaged God. We have an intellect, it's what keeps us different from the animals. This was, of course, the first way we attempted to build

cognitive AI, to see if we could isolate our reasoning process. When we get to expert systems, when we get to the tools that we have in AI, we see AI as mimicking our functions. We also saw, secondarily, that we are called to function like God in that call to have dominion over the earth in the Bible.

Noreen Herzfeld:

More recent theologians, particularly Biblical scholars, said, “This is a better way of looking at the image. It’s more dynamic, it fits better with the text.” Fine, let’s fit that with AI. If we are God’s hands on earth, AI could be our hands in places we don’t like or want to go, like Mars. Here we have the traditional idea of technology as a tool. It extends ourselves, and the question, of course, that this whole day has been asking is, how far can we make that extension go? If we have a tool like a hammer, it’s an extension of our arm, but you still need the arm. If we have a tool like the Mars rover, it’s an extension of our presence, but we’re not actually holding onto it. We’re not actually present ourselves with it, except over the internet.

The third way we’ve looked at the image of God is as relationship. This is a particularly Christian interpretation of the Genesis story. Karl Barth is probably the biggest exemplar of it. He said, “The Christian God is a Trinity—Father, Son and Holy Spirit—always in relationship.” Therefore, God is a relationship. We image God, not individually, but in our relationships with each other. Here again I think this reflects what David just said, about our necessarily being social beings.

Noreen Herzfeld: Of course, we look at AI, and we have to ask the question, do we want to step away from the functional view of AI just as a tool, as an extension of ourselves, to a more relational view of AI as a partner? Many of you have mentioned that we're already starting to make those first steps. We see these steps, of course, in our stories, but we also see them with Siri. We see it in robotics, we see it in the development of sex bots, where people are looking for a very intimate relationship with our AI. In many ways it is the relational angle that we're beginning to look at.

The question that I ask is, how real is this relationship? Here are four criteria that the theologian Karl Barth gives for what is an authentic relationship. To look the other in the eye, to speak to and hear the other, to aid the other and to do it gladly. Let's take a real quick look at each of these.

To look the other in the eye. Here's a cartoon, one of my favorite cartoons from the *New Yorker*, where one dog says to the other, "On the internet nobody knows you're a dog." We know that technology can actually stand in the way of looking the other in the eye. We've heard already from a couple of our speakers about how cellphone technology seems to be breaking down certain social abilities for the younger generation. We also have people in robotics working very hard to give us computers that we actually can look in the eye, that have a definitive physical presence.

Noreen Herzfeld: How about speak to and hear the other? Well, here we would think that computers are probably the strongest. We speak to them, they

hear us. They speak back to us in ways they never used to before. As other speakers today have already reminded us, there is more to speaking and listening than just words. I think the very use of emojis tells us that we know this. We know the words by themselves are not quite enough.

Noreen Herzfeld:

Can we aid the other? Sure. Computers are aiding human beings in multiple, multiple ways. What about do it gladly? Can the computer do it gladly the way another person can? Here we get into the realm of emotion. We have to ask the question, does emotion need a body? If we look at the four stages of emotion, and here I'm using the stages as put forward by Jerome Kagan, it's perception of a stimulus, a change in feeling that is sensory, appraisal and response. Clearly AI can do the first quite well, perceive a stimulus. Clearly it can appraise that stimulus and clearly it can respond, in some way, to that stimulus.

The real question is number two, what about a change in feeling that is sensory? In other words, unlike Bill Clinton, can a computer really feel your pain? When we think about empathy, for example, as an emotion, you perceive someone else in some sort of difficulty. If you jump straight to number three, and you say, "Okay, let's appraise this. What should I do to look good? Now I'm going to respond," but you miss step number two, well, we have people who are like that. We tend to call them sociopaths. It raises the question, would a relationship with a computer that does not have that step two change in feeling that it's sensory ultimately be as unsatisfactory as a relationship with a sociopath?

Noreen Herzfeld: I am going to end here with just a little more quick theology. I believe that the Fathers of the early church actually recognized the need for our physical bodies. Going back to something Marc said in the previous panel, he said, “Well, should the computer be our savior?” There certainly are people who are looking towards AI, either to be our mind children or to be a place where we can upload our mind, something like that. The Church Fathers knew that without a human body, a human-like body, without one particular locus, we would have nothing like a human-like reality.

The Christian church also recognized this and said, “We believe that God came and shared our bodily situation with us. That is how we know that our God has empathy, from having had a human-like body.” I’m going to end on that note, because Ben is telling me I’m out of time.

Ben Johnson: Very empathetically I might say.

Noreen Herzfeld: Thank you. You felt my pain.

Ben Johnson: I did.

Amy Pritchett: Well after these two speakers, I’m feeling very humbled. I’m afraid I’ll be far more pragmatic. I find great joy in flying in a cockpit. I don’t mind driving on a windy mountain road either. My own specialty is studying how experts interact with technology.

Amy Pritchett: With this I’ll speak much more like an engineer, and talk about some very pragmatic things that have changed the experience of experts, because of how they interact with these intelligent machines. Given the forum today, I

thought I would talk a lot about some of the more ethical challenges that we have, the legal challenges, and the broader framework of what happens when we put these experts into systems.

Amy Pritchett:

Now I'll start with something a little more general. We talked a lot about self-driving cars, this was a focus our speakers earlier today. I'll point out the juxtaposition here, this was the first Tesla model's last crash, the picture on the lower right. Surrounding it I have the corresponding page that was pointed to in the accident report from the driver operations manual. Pointing out that there's many factors that can impact the performance of this automated system, including poor visibility due to heavy rain, snow, fog, bright light, etc. Damage or obstructions caused by mud, ice or snow. Narrow or windy roads, and interference from anything outside that might impact ultrasonic waves. Warning, the list above does not represent an exhaustive list of situations that may interfere with proper operation. Never depend on these components to keep you safe. It is the driver's responsibility to stay alert, drive safely and be in control of the vehicle at all times.

Amy Pritchett:

What are we finding? We are putting human experts or in the case of a driver at least somebody who's got a driver's license in a context where we go, "Hey, if you believe the Elon Musk tweets, hey it's self-driving, yeah." If you read the fine print, it's not. Now the real love of my life is in air transportation. An air transportation is operating in a context where the Federal Aviation Regulations part 91,

operations of aircraft starts with after the definitions.

Amy Pritchett:

The responsibility and authority of the pilot in command. The pilot in command of an aircraft is directly responsible for and the final authority to the operation of that aircraft. In an inflight emergency requiring immediate action, the pilot in command may deviate from any part of this rule to the extent required to meet that emergency.

The two parts here, create a very intense human reaction. I love working with people like the Air Line Pilots Association, because many pilots deeply internalize the notion of this command. That comes from also the captain at sea. Needing to be responsible for the outcome, and being given the authority to do whatever it takes, including bending any other rule to keep everyone safe. In this context then, the pilot, the human is very much responsible for the final outcome, even as they're increasingly being given machines that are not always perfect in all situations.

Amy Pritchett:

Even twenty years ago, we defined a paradigm that said "the pilot is the final responsibility." This cartoon is actually twenty years old. That is in a current air transport aircraft, including the 757 and 767 that first flew in 1979. There's not much, in the ideal situation, that the pilot needs to do. After configuring the systems, and loading the flight plan while sitting at the gate, the pilot needs to turn on the engines. There are two pilots. The captain on the left is the only one with the tiller wheel to steer while taxiing, so the other pilot needs to turn on the engine. The captain needs to taxi out and line

you up with the runway. Then either pilot—the captain or the first officer—needs to hit a button labeled “take off, go around” (TOGA). The airplane will then take itself off. The pilot just needs to raise the flap and gear. It’s not worth automating. Then, sixteen hours later, the pilot needs to put the flap and gear down. The airplane will land itself and stop on the runway, at which point you say, “Thank you auto-flight system.” The captain has to taxi the airplane in and turn off the engines. That’s the theory.

Amy Pritchett:

Now, of course with this we shouldn’t be surprised that pilots report hand flying the airplane on average somewhere between two to seven minutes a flight, even on the longest flights across large sections of ocean. Only two to seven minutes of flight. By the way, those two to seven minutes of flight are often the takeoff at an airport that does not have the ground facilities to support an automated takeoff, or landing at an airport that does not have the ground facilities to put on automated landing. The case of a sudden air traffic construction, asking for an immediate climb to center of return and any emergency. Other than that, we don’t need the pilot.

Amy Pritchett:

The staff in model that we now need to think of for the pilot, is not a measure of their continuous work load in flight. It’s more like staffing a firehouse, thinking about, how bad can the emergency get and when it gets bad, how many pilots do I need to have there? Likewise, the pilot’s job now is one of a supervisor: sitting there, monitoring. “Why is it doing that? Is it doing the right thing? That doesn’t look right to me. Should I intervene?”

Oh no, that one went okay. Oh no, that didn't go okay."

Amy Pritchett:

This is a very hard paradigm I think for a modern job. Indeed now that increasingly I'm a supervisor, I find it's actually a very hard job. What we are doing is we are pushing all sorts of people who are selected for manual control skills, for mechanical ability, and we are pushing them to become supervisors as part of their everyday experience.

Now, this is something that we've been studying in commercial air transport. This is from a 2013 report, a broad based public private partnership. It includes the performance-based aviation rule making committee headed by the FAA. Then looking at what is the role of the human in such an automated context, their first finding was that, the entire aviation system is predicated upon the belief that pilots will step in and manage risk. Foresee problems coming, prevent the problems from occurring in the first place or responding to them when they show up.

As part of this report, we collected a significant amount of data that'll step through here. In this case we looked at what's the frequency of aircraft malfunctions. The definition here was anything requiring pilot intervention for the flight to continue normally as originally intended to the intended destination. These malfunctions may range from the very bad, engine on fire, emergency landing at Philly, that we saw a year ago to a circuit break or popping. Just need to recycle the circuit breaker from the fair like benign but still requiring some human interaction, to the very severe.

Amy Pritchett: Not surprisingly in red we see that accidents involve malfunctions on the airplane, about fifty-five percent of the time. Likewise, major things that qualify as major incidents, or things that are reported to the Aviation Safety Reporting System as incidents have an incident. What amazed us was the green, these were from Line Operation Safety Audits or LOSA, check ride pilots riding with pilots. Riding in the cockpit, surveying how's the operation going, how are the pilots doing? The number here, something like ten thousand flights across all operators in America, all fleets chosen to be a broad sample. In that case, twenty percent of the flights required pilot intervention. Again, it may just be crunchiness circuit breaker.

To me this was amazing. Before we had not captured the level of pilot intervention. Indeed, it had been common to hear that pilot air was implicated in ninety percent of accidents. Bad pilot, let's automate them out. When this data came out, people said, "Why didn't we know that 20% of the flights that pilots have to intervene?" The answer is, it's so common that the pilots think it is part of their job. They don't get off on in five flights and go and file paper work saying, "I had to push a circuit breaker today." That's just part of their job that is unremarkable.

Amy Pritchett: Indeed, if we look at the human contribution in current day operations, in 20% of flights, something technical fails and the pilot resolves it. In 55% of accidents, something technical fails and the pilot can successfully resolve the failure or control the situation. Indeed, a significant portion of accidents also involve

pilot error. Often these days the pilot error is listed as over reliance on a machine or poor monitoring of a machine.

Amy Pritchett:

When we talk about the human or the machine, which is safer? Please just consider that these are numerators or conditional probabilities that need to be multiplied through by a large number in the upper case. Twenty percent of flights is our large number, please nod, yes, that is a huge number. Fifty-five percent of accidents is, fortunately, a very small number. When we look at the human contribution of safety, we find that there's still an incredible amount of human contribution that we need to consider.

Some other quick things, same study looked at how often is there a flight management system program in there? It's huge, it's about 26% of flights. There's some difficulty by the pilot in interacting with the machine at the level of entering in a command into the flight management system. This happens across accidents, incidents, and normal flights. The strongest predictor of an accident that we could find was a breakdown between the captain and the first officer in their communication. There's some secret sauce that we do not yet understand well enough to then put into the machine about two humans cross checking each other.

Amy Pritchett:

With this, my final thoughts here, is that machine autonomy, we think this is novel because machine stepping into what humans have done. This is an agency we give to the machine via things that have been created by humans. Datasets created by humans or

engineers created by humans. This is an interesting thing, but is it necessarily good? Should we really be focusing so much on replacing human work? Twenty years ago, as these automated cockpits were being brought in they were saying, “Oh pilots will love it. It will make their job easy. It will reduce their workload.”

Amy Pritchett: I think also in a human spirit, it’s important to think about for many pilots they describe being a pilot as their calling, their vocation. This intrusion into the notion saying that, “Oh we’re just helping people with autonomy,” is not always as directly simple as we have attributed it to. Thank you.

Matthias Scheutz: Alright, super, thank you. I just learned something new today, that you can actually give a presentation via Zoom without needing the cable. For all of you who’ve struggled having the right adapter, that’s the solution.

Matthias Scheutz: Alright, Ben asked us, as David said, some good questions, in particular how autonomous machines could affect us as humans? How they’re affecting us now and how they will affect us? That’s what I want to talk about. I want to give you a dialectic, an argument in both directions. I’m starting with premises that are based on just working human robot interaction. I do both AI and robotics type work, where we develop technologies for intelligent robots, autonomous intelligent robots. Then we also evaluate them in human-computer, human-robot interaction studies. There’s a whole research community that also does that empirical evaluation, and that’s why we know how people in some cases perceive

intelligent systems and autonomous systems and what that does.

Matthias Scheutz: What I'm listening here is a summary of converging evidence for it. Autonomy increases the perception of agency in humans. Agency perceptions increase mental state descriptions that you believe that there is somebody at home in that agent. Both of those happen automatically, you don't have to think about it, it just happens. You see the Roomba drive around, and you think it has a goal to go somewhere when it really doesn't.

Human likeness is a modulator, the more a system looks like a human, the stronger the effects typically. It's easy to perceive autonomy even in simple machines, like the Roomba. You look at it, you can see it is self-propelled, it moves on its own, it seems to have a purpose. We do that, and I could tell you a long, evolutionary story about why that is. What I'm going to do is this, given that, you can say that these mental state inscriptions can help make sense out of autonomous system's behavior. It's easy to instill them in us, a system can communicate what it's doing or what it's closer that way. That could lead towards positive attitudes towards the machine prefer and also ultimately trust.

Matthias Scheutz: On the other side, the mental status inscriptions can lead to unwarranted emotional experiences. For example, people might feel gratitude, which they do with the Roombas for cleaning. Or they may feel guilt or jealousy or adoration and others. Feelings of acceptance can lead to increased purpose and sustain human dignity. You think the machine likes

you, accepts you, and as a result you feel good about it. Or, the machine doesn't act in a way that you understand, it's not a chat bot, it's a socially assistive robot that helps you out of bed. You feel gratitude for it, you want to talk to it, but it's not made for chit chat. Now you feel rejected, you feel it's ignoring you. It can lead ultimately to a loss of purpose and dignity.

Matthias Scheutz:

Some of the consequences of interactions with autonomy might be beneficial to humans. Give you a feeling of empowerment and social inclusion, and may be a gain of dignity. Conversely, you get interactions that may be detrimental, and you feel inferior to the machine, you might actually lose social contact and ultimately loss of dignity. While these human social interactions with autonomous machines could improve human interactions and in general human societies, they may do the opposite as well.

There's, perception of autonomy is something that has been studied for a long time from the early 2000s on. There's lots of evidence for it, how people perceive autonomy in machines, whether the autonomy is there or not. Often times the studies are done with remotely operated robots, people just not told that they're remotely controlled. Agency same thing, there's a lot of literature on it that tells you that you very quickly perceive agency in these systems. That goes back to the psychology work and geometric shapes moving around where people interpret intentions and feelings and everything else into those.

Matthias Scheutz: Similarly, human like appearance, there's a lot of work on what that does and how that triggers theory of mind areas in people. In fact, there are MRI studies that show depending on the appearance, you will have different activation of the theory of mind areas in the human brain. The same goes for human like behavior, in particular language. If the system interacts with you in natural language, that is one of the strongest indicators of human like behavior. It comes with a host of inferences that people make automatically as a result. I'm not going to go into any details, but I'm happy to talk about it afterwards when I go through a list for discussion of beneficial and detrimental effects.

One thing they can do is increase the focus of attention. You're losing your attention, the system comes back to you and says, "No, this is what you're supposed to do, keep working at it." It might actually engage you. Could help you understand emotions. We're already using robots for kids with autism, we're talking about sociopaths. You might actually by interacting with an autonomous machine, learn a little bit about emotional states and they're used that way.

Matthias Scheutz: We can use the machines to prevent harm. The big argument for autonomous cars is, they're better than human drivers. Arguable point, but that's companion robot, same thing. Might prevent psychological harm. If you have an aggressive child and the aggressive child needs to get rid of the aggression, well maybe let the child interact with the robot and get it out of your system. People have argued that. Social engagement, if you have difficult engaging with

people, maybe you start with a robot, which is a safe space and a safe place to improve these skills.

Matthias Scheutz: By interacting with machines, you might personally feel empowered, because the machine takes you seriously. It gives you the attention you need, gives you the emotional support, might improve the mood. We see that in care domains for example, with robots like the PARO, which is a robotic seal that people stroke and pet and then they feel better about themselves.

It might preserve human dignity. If you cannot eat by yourself and you have to constantly ask other people, because you're quadriplegic to help you feed yourself, that's a loss of dignity. As opposed to a robot that just does it for you, and you can now go with that robot to Starbucks and order coffee on your own. Overall, these robots could make up for difficulties we have interacting with other social groups, and therefore strengthen social interactions and human societies.

Matthias Scheutz: Here's a list of seven points and I'm going to have for each of the seven, the seven opposite. Distractions, think of cellphones, massively distracting. These systems could be distracting too. In fact, we see when Roomba vacuum cleaners vacuum, people watch them and don't do anything else. It's like why did you buy it in the first place? There could be emotional and social harming. I think Allan already mentioned that there's a literature on it, that you get when you play too many of those first-person shooter games.

Matthias Scheutz: There are social changes in norms for example, I think Amazon had to put in please before commands, because the kids listening to how you talk to Alexa made them talk that way to other kids. There's this transfer effect. I've written 10 years ago about the danger of unidirectional emotional bonds that you might have with robots. You are invested in the system, but the system cannot reciprocate. You are grateful that it does something for you, like get you dressed in the morning, get you out of bed. You want it to like you, but it doesn't have the capacity for it.

Autonomous cars is another example for harm. Social isolation, we're already seeing that with cellphones and social media, that people prefer to interact with other people in a mediated fashion through social media. Now imagine you come home and there's your robot, and it gives you everything you want. In fact, you mentioned sex robots, we wrote about that too. That might be a very good example of how something that's intended for a particular purpose, may be broaden into intimacy and a real relationship rather than anything else. We see that already with people who don't use robots yet, but as you said they're being produced and that's a particularly bad case I think. It's great danger there.

Matthias Scheutz: There's another interesting dimension, which is something that Austrian philosopher, Gunther Anders, who went to the US in the 50s wrote about. That's the human perception of perfection in technology. You look at these machines that are so powerful and so fast and can do things so beautifully, and you feel inferior automatically. You feel the awe and he

called that the Promethean Shame. If you look at these robots what they might be able to do, that's the experience you might have.

Matthias Scheutz: We might get to the point where machines are going to tell us what to do. It might happen in very gradual fashion, from the pat down at the TSA, because people will prefer robots to do that instead of people based on studies. To various other roles that robots may assume, where they are in a position where they will tell you something. Eventually by having these machines in our households and by us preferring to interact with machines rather than with people, which could have a detrimental effect on human societies at large. There you have it, I'm done.

Ben Johnson: Wow, that was impressive. We have time now for questions. Noreen, I have a question I have been wanting to ask you. I'm a Tolkien nerd, I named all three of my children after Lord of the Rings. Oh, this is going to be recorded, great.

Ben Johnson: Tolkien is this notion of sub-creation that is an act of worship in some way. This is how you described AI in some way as humans doing as God did in creating us. We are creating AI and passing that forward, but if you turn your pages forward in Genesis a few bits, you get to the Tower of Babel. Where a group of people decided that we don't really need God and we can do all of this, fine, look at this money as cheaper or something, that ends very, very badly. I'm curious now, how do you think through as a theologian in some sense and you observe this and you read Genesis? What separates, in some way, this good thing of

doing as God has done in acting in the image of the creator versus trying to subvert it or overcome it in some ways? Is it just a matter of motive or does it actually play out in some way that we can observe and can it guide in how we deal with autonomous systems going forward?

Noreen Herzfeld:

Yeah, that's an interesting question. That was actually the slide I skipped—unintended consequences. If you read Genesis four through ten, you actually get a very succinct description of the development of early technologies. First the development of agriculture, and then you get Cain and Abel. That leads to murder or what signifies, which is the fights between nomadic peoples and settled agricultural peoples.

You get to the Tower of Babel, which represents the whole development of cities and the city state, and the problems that that brings about in separating people, warfare, all of that stuff. Finally, you get Noah and the flood, and the ark which is a fairly good technology for saving some stuff. You do not get, until the story of Noah, the first covenant between humans and God. I think what all the stories are telling us is that after humans get created and then we become creators, when we think we can do that entirely on our own we tend to screw up. That it isn't until you get to the Noah story that you get a description of a technology in which we are, in a sense, now in a covenant with God, trying to do it with God.

Noreen Herzfeld:

I think if you extend that thinking to AI, it repeats something that we've heard over and over again today, which is that composite systems, systems where humans and

computers cooperate and work together, tend to work the best.

Ben Johnson: If you have questions please come to the microphone. Do you panelists have questions for each other?

Amy Pritchett: I was just thinking of Jacob wrestling with God too, the same way we sometimes wrestle with our machines.

Noreen Herzfeld: Yes, we do.

Audience: Hi, I'd like to come to the topic of emotions and the role they play in the interaction and relationships, and person to person as well as person to machine. Our limbic system evolved as a precognitive system, which controls a lot of different aspects. Motor movement and the expression of emotion also. It's not something we can control, at least not very well. It reveals something of our inner state, and it's telling the truth about our inner state, which helps with the relationship building and the reciprocal trust that has to happen in any kind of relationship.

Audience: Do our robots need to express emotion? I know there are robots for example, the baby harp seals, they wiggle and they express comfort when they're being petted. It helps the senior citizens in some way, but it's lying. They don't feel anything. I'm having a problem with that, because it's not an honest signal about the inner state of the robot. I really want to know if I can trust my robot, but if it's faking emotions, overconfidence, "Yeah, Dave I can take on that task." When it's really going, "Hmm I don't know." Do we need to have these authentic relationships, do we need to

build in authentic emotions that reveal something of the inner state of the robots so we can determine whether to trust them or not?

Noreen Herzfeld: Yeah, I think you're onto something there. That that is a problem. That right now what we have is the two ends, but the middle is empty. The robot can sense our emotions, it can make a calculated response, but we know there's really "no 'there' there."

David Danks: Yeah, I don't want a sociopathic robot.

Noreen Herzfeld: Right. And in some way, this, I think, mimics a difficulty that came up earlier today with so many machine learning systems, is that, we give it the training, we see the outcome, but we actually don't know what it learned. We don't know what that outcome signifies and that makes it very hard to trust the system. You realize at some level that while you might think it has learned how to identify faces, maybe it learned how to identify cloudy days, and that it can suddenly fail. I think that that black box middle is a problem for trust.

Amy Pritchett: Go on David.

Matthias Scheutz: I wrote a paper a few years ago that I called the affect dilemma, and that's exactly that question. The dilemma is like this. When we don't build machines that have genuine emotions, and we can argue whether the body needs to be part of it in the human like sense, emotion theorists have different views on this. If they don't have genuine emotions, whatever those are, they can fake it and they can manipulate us. We don't want that. They give us signals, but you're right, it's pretense and that's what almost all the

social robots are doing right now actually. Hold on.

Matthias Scheutz: If you do the opposite, if you put machines that have genuine emotions, then you're creating machines that can suffer. There's a German philosopher, Thomas Metzinger who has specifically argued that we're already doing this, we just don't know it. We're replicating the causal structure of what it takes for some systems to instantiate these emotions. You're caught between a rock and a hard place. Do we want to create robots that have the capacity for suffering and pain?

David Danks: I think the emotions in a functional sense and the suffering of a robot doesn't really have a function unless it's saying it's fallen down, it's got a broken leg, "Help me up."

Matthias Scheutz: We're going to get into philosophical discussions.

David Danks: No, it's not necessarily huge, it's not bitter sweet.

If it's about a robot's inner state, then I think it's an open empirical question we're not able to answer. I don't actually think anyone really knows. If it's not about inner state, then does it matter whether it's done with emotion or not? The problem with the sociopath is in some sense often not that they don't feel emotion, it's that the emotion's not gotten in the behavior in the right way. They know how to fake the signals to deceive everyone around them.

David Danks: Suppose you had a truthful, earnest robot that didn't feel emotions, but could reveal its inner

state. For trust we do need to know something about the inner state of the robot, that's what enables us to have the ability to predict how it's going to behave in new circumstances, and to generalize from single cases. I think we need to keep separate "emotions as a particular evolved mechanism we have for signaling inner state" from "the need for inner state information." We might imagine having robots that just reveal their inner state. They wear their heart on the sleeve as it were, and don't have to have emotions. We use emotions because it's a way of conveying really rich inner state information in a very compact signal.

Ben Johnson: I'm going to let Amy jump in here though.

Amy Pritchett: Yeah, and I'm sorry, I'm going to be so pragmatic. I'm in amongst the philosophers and the theologians and I'm going to say, it's a designed system. What you're really interacting with is another human's vision that they have instantiated in a computer and put into place. If it wiggles and looks cute, then that other human was successful in making something that would touch you. I think there's a parallel there with an author being able to touch you through their communication that they put on the page and then it took on its own life as it went out into society.

It's not human and machine interacting, it's human and human interacting through one person's conception that they put out into this machine that goes out and interacts with others.

Audience: Thank you.

Noreen Herzfeld: To come back to what Matthias said, we don't perceive that. We give the agency to the machine, we attribute what the machine does to the machine. While I think you're absolutely right, most people don't know that.

Matthias Scheutz: There's evidence for it. We've done experiments where we tell people explicitly that the robot is teleoperated, so that when they're talking to it, they're talking to a human on the other side. Yet they talk to it exactly the way they would talk to it when it was a robot, an autonomous robot. Somewhere it's lost that there's a person, when there's a body in front of you that interacting with.

Then very much to David's point, I think the reason why so many social robots fail is because they don't actually have a corresponding internal state to the particular kind of expression that they show. It's not systematic in the way it is with us. That means they will smile, because you're smiling, but you just talked about the death of a family member and it's an embarrassed, a hysterical smile. It will smile back at you because that's what it does in that case [crosstalk 00:58:26]. Exactly, yeah.

Audience: I'm looking at the title autonomous systems and the meaning of humanity. For many of us we find a lot of meaning in life through our occupation, through our employment, through the things that we do. I'm thinking back to a conversation I had with a person that when I was in Thailand, who I met, his job was he was a human de-miner. He would take a knife and crawl along on his hands and knees, and stick

the knife in the ground trying to find live mines in Vietnam and Cambodia.

Audience:

When he told me this I said, “Well, that’s a great job for a robot, we’re trying to build robots to do these types of things.” He got very upset, he said, “Why? This is my job, I’m very good at my job. This is not a risk for me, this is how I feed my family.” How do you think autonomous systems will impact employment, employability in the next 10, 15, 20 years? How will that affect humanity?

David Danks:

I’m not an economist, I’m not going to pretend to answer that question. I did though want to connect it to something that Amy said that I think was really important, and I wanted to make sure to highlight. Which is, the fact that when we think about impacts on the workforce, we have to think not just about economic impacts and this is connected to the story you’re telling. For those of you who are in the room who are professors, if there were a robot that could do your job, would you only be out of paycheck? I contend that for most of us that wouldn’t be the case. Most of us would suffer very real psychological and moral harm, because something that’s critical to our self-identity has been taken away from us.

David Danks:

You mentioned with pilots and I think to tie to the meaning of humanity part, I think that a lot of the analysis of workforce implications focus solely on the dollar values or the displacements as a number. I think we need to be talking a lot more about the harms that are done to people when they lose jobs, not just in the paycheck, but in the sense of ripping apart some of their

identity in a lot of the changes that are happening in more and more workforces.

Amy Pritchett:

Thank you and indeed I glossed over a middle point in my last slide before Ben tackled me, but we have to understand what the work really is before we try to replace it. An example I love is my twelve-year-old girl coming home on a hot day and telling me about how a bee got on the bus, the windows were open on the school bus. A wasp or a bee got on the bus, can you imagine what happens in a middle school when there's a wasp? There are kids yelling, jumping over seats and there's mass panic. The school bus driver had to pull over the school bus, pick up some piece of paper, stomp back, kill the bee, get the kids settled down and then keep driving.

It's common when we talk about self-driving cars or school buses to say, "Oh we're replacing the driver." That made me reflect that the school bus driver is taking fiduciary responsibility for those kids. The school bus driver's trained on first aid. The school bus driver's a disciplinarian who's breaking up bullying. If we replace the driving function, we would probably decide we needed a teacher to ride on that bus to handle all those other human functions.

Amy Pritchett:

Likewise when we talk about a totally automated car, I personally love driving, I wouldn't use it. I would love it to carry my kids, but if it carries my kids, it has to take custodial responsibility. When it drops them off it has to make certain that they are with the people who they are supposed to be with. If there's a wrong address it needs to be able to recognize it's

dropping them at a crack house and not at the best friend's. It needs to help me resolve that situation. We do not properly associate that even the most mundane jobs, what we think as the most mundane jobs, have many layers of interaction that would need to be captured if we're really going to let a machine do it without also needing a human there for the other aspects of the work that they do.

Ben Johnson:

David can I follow up on this and though I ask, there's a classist angle to this. We've been replacing low skill blue collar jobs with machines for a very, very long time. We did not have big conferences, we did have a communist revolution. We didn't have big conferences about this type of concern. How much is it really is the fact there is the, back in the last presidential campaign there was this Ted Cruz this commercial in the primary with all these lawyers, investment bankers running across the border trying to signal, "They're coming for your jobs too," so the public should be very worried.

How much of this is actually new and how much of it is just hitting a new group of people for the first time and now people are freaking out?

David Danks:

I think that's a good question. I think it's mixed. I mean of course the Luddite rebellion was exactly a lower-class profession being automated away and fighting back against that. I think there are certainly a classist element to it, the kinds of folks who come to conferences like this are the ones who are feeling threatened. I actually think that there's a deeper

issue that goes to a point that Noreen was making about the role of reason.

David Danks:

Since well prior to, but Aristotle most famously said “man is the rational animal.” It has always been that which was held us up as why we’re not like all the other creatures in the world is our rationality and reasoning, it’s our ability to thereby use language. That’s the thing that’s now being threatened. We’re not just using machines to replace human physical labor, we’ve always done that since we domesticated animals. We’ve been replacing human physical labor with something else’s labor. Now we’re replacing human cognitive labor and I think that actually threatens another thing that’s at the core of our identity, which is our rationality in reason is, that’s the thing humans have that nothing else in this world has.

I think that’s another reason that we feel this threat in a different way, is that the farmer to a certain extent, I’ll speak to my own grandfather, didn’t mind when he got a tractor. He didn’t have to be out in the fields quite as much and didn’t have to walk behind the ox or cattle in Southern Illinois. Now we are that ox, we’re the thing getting replaced. What has happened to us? I think that is an important distinction.

Audience:

Thank you. As we think about humanity, I’m interested in how when we think about some of the stuff that’s going on in Silicon Valley in terms of who’s creating these machines, largely white, largely male, largely affluent. That there are a lot of situations where have shown that actually women have been forgotten about. Used to think like excuse me, airbags are a

famous example. They were designed to account for a standard male, they didn't realize that actually women would be in cars or just didn't particularly care. Or with things with face recognition where they didn't use any datasets that included women, darker skin people, and younger people. When we're talking about humanity, can you talk about who's humanity as well or when you look at immigration or things like that, it may just be a certain group of people's humanity, as well.

Noreen Herzfeld:

I think this has already come up during the day when we've seen that it's so easy with machine learning to put our biases into the machine just by the data sets that we use. I've been struck by the ideas, for example, that you get people like Ray Kurzweil, who thinks we could upload our brains to computers. Then he'll say, "Oh, that doesn't mean we're getting rid of sex in case you were worried about that, because sex is all in the brain anyway. We could still have this experience." Yet we have to say, "Well wait a minute, what kind of an experience is that?" It's a totally by yourself—masturbatory—in a sense experience.

The interesting thing I found is that these are all men who are suggesting that we could do without a body. That we could be just an isolated mind—"Oh, but we can still have sex." I just never heard a woman who would think that way.

Matthias Scheutz:

That's a good one. Let me say one other thing that I found striking and it's now fairly well established. People transfer attitudes they have and experiences they have from humans to machines, often that goes automatically. Sex

and gender being one of them, and it goes with all the effect. One of the things to your point, that I find very striking is that stereotypical sex expectations or gender expectations or biases transfer to robots directly. That is shocking because it doesn't take very much to make a machine appear gendered, just give it a voice. If it has a male voice, it's a guy. If you put red lipstick on the lips, it's a female. I mean there are experiments like this.

Matthias Scheutz:

The moment you do that, you imported the host of prejudices and stereotypes that people have. The robot will be construed that way and the interactions will unfold very differently as a result of it. That's something that actually, you talked about the design, that's where the design decisions are so critical. That we be really, very cautious about what we put on a robot and how we design. The Pepper, I don't know if you've seen this, looks like it's wearing an apron. Most people will say it looks like a female.

The Nao robot is short, everybody will say it's a kid. You give it a deep voice and people go, "What is that?" It's that short. It shouldn't matter, it's a machine. You should be able to combine that any way you want, but it does matter to us.

Ben Johnson:

Two questions, we have five minutes and two questions I would like to get to.

Audience:

Okay, I'll make it quick. Thanks, great, a set of tops, thanks for a lovely panel. My question is for Noreen. It seems like the theology that we spent most of our time talking about was Judea Christian theology, which the southern Baptist, Amy likes and my better half who is Jewish

would also appreciate. I wonder if we also might gain some insight into other nation states and cultures might regulate the deployment of autonomy in their societies by thinking about Hindu perspectives from the Rig Veda, Buddhist perspectives, Confucian perspectives especially given the graph that we saw from Paul last night that had most of the UAVs being deployed in the world coming out of China. If there's any insight there and what it would look like.

Noreen Herzfeld: Well, there certainly is and there just wasn't time to cover the whole spectrum in twelve minutes. We can clearly see in Japan that robots are much more accepted—in the work setting as receptionists, in a home setting. A new robot in a work setting will generally be welcomed with a Shinto ceremony, because in Shintoism you have a form of animism which believes that everything is ensouled. Even inanimate objects. I mean, even this watch has a certain amount of soul. Then imagining working with a robot, having a religious ceremony with a robot, this is not an issue. The robot can be seen as ensouled in a way that we would not, in our Judea Christian tradition, tend to think.

Noreen Herzfeld: You also find, for example, a video out on YouTube that you can look up. It's Pepper, but instead of wearing an apron, Pepper is wearing the robes of a Buddhist priest and is reciting the prayers for a Buddhist funeral. When you think about different forms of Buddhism, for example, in Tibetan Buddhism, the whole use of prayer flags—the idea that if you put your prayer on a flag and put it out and the wind blows it, and the prayer is, in a sense, being said. Then having a robot that repeats the

prayers can be seen as having the same function. Again, that's a different understanding of prayer than we would have in our tradition.

Ben Johnson: That's great.

Audience: My comments are related to the theology and robotics comparison too. With your examples, the Genesis in the Bible was written some 2,000 years ago and has been changed a little bit, but stays more or less similar. That was if you look at some of the stories and so on, there is this almost a theoretical conception of creation and what we think creation is like in the stories. What you said was that there's some similarity to what we are doing today in implementing some of them to the extent that we couldn't have implemented them.

It's no surprise that our story then and the story now has similarities. Are there differences that we did not see when we wrote the stories of the Bible and when we're implementing AI today that are there star differences that you see?

Noreen Herzfeld: Well, I was looking for the similarities there in seeing that these stories capture something deep about the meaning of humanity, the title of our talk. Of course, there are differences, and this would be a whole different talk going in a whole different direction. When you have a People of a Book, as in the three monotheisms of Judaism, Christianity and Islam, all have a text. As you point out, that text was written a very long time ago and in a very different cultural context.

- Noreen Herzfeld: One of the problems that we have is that, we reach a point where it's difficult to understand those stories because we no longer understand the cultural context in which they were written. Also, the stories lose a lot of their purpose. What I mean here, I call it the "aha factor." I believe that many of the stories that are written in these texts are actually meant to hit you upside the head and make you perceive the world in a different way. When the stories no longer come out of the culture that you live in, it's very difficult for them to do that, because you have to explain the story. What's kind of like explaining a joke. If you tell a joke to somebody and they don't get it, and then you explain it, at a certain point they're going to say, "Oh I guess I get it," but they're not going to laugh. It's the same thing with these texts. As our culture moves away from these texts, we need to explain what is trying to be taught by the story being told, but we lose the "aha factor" that hits you upside the head and makes you see the world in a different way.
- Ben Johnson: Well, thank you panel very much, thank you all for being here. Give the panelists a round of applause. I believe Admiral Houck has final parting words for us?
- James W. Houck: I do, and so this is the point where you're waiting for me to say, "Thank you," and you're dismissed. Just a few words before we get to the dismissed part. First of all, thank you to all of you who, and I'll echo the words earlier, even though now it's only 3:30 for those of you who have remained throughout the day. I hope you found it rewarding to sit through the whole spectrum of the program that we had today.

James W. Houck: Since Noreen Herzfeld outed herself by copping to be the person that made the comment about feeling uncomfortable, I'll out myself and say that I specifically thought that this particular panel was an important way to end the day. I think it's easy—I'm an expert in none of this—but I think that it's easy for us to get wrapped up in the component parts of this discussion, in the technology of the discussion or the law of the discussion or what the what have you of the discussion. I think this panel was a way to lift us up above that a little bit.

I was really struck by Marc's really excellent presentation and Marc's comment. I'm going to make an assumption about Marc's age and my age. I know what my age is, I'm going to assume that Marc is a good bit younger—a lot younger—than I am. It's common that his generation, if you will, is looking for convenience. It struck me in the conversation about, there's so many echoes of different parts I'm going to try to integrate with this comment. The comment that Claire made about losing some of our skill or losing some of our functionality. The comment about the person who is disabled and would want a robot arm to help them in a way that will enable them.

James W. Houck: In my own reaction, and I have a point to this, I'm getting there. My own reaction is that there is some value in suffering. That there's some value in hard things. That the search for convenience, here I wonder is it the wrong search sometimes and are we losing something by looking to make things easier for ourselves? I'll admit, my kids will hear me talk this way and just say, "You're old, just go away." Those are the kind of interesting conversations and I

think that all of our panels, I hope, I know, were bringing out.

James W. Houck: This was a hard event to put together. I mentioned to somebody last night, that for me, it would have been pretty easy in about a week or less, I could have come up with a room full of lawyers to sit down and talk about these things. That's not what we've got, and I think this has been challenging on a couple of levels. First of all, it's challenging to get scholars and academics to talk across their boundaries.

I know I've been here for seven years now. I know enough of them to know that this can be a hard thing to do to get people to do that within the academic community. That's one dimension of the hardness here. There's another dimension to the hardness, which is to include in this discussion not only scholars, but the public and students. I know that all of you are represented here today in some measure. On a couple of different dimensions, we're trying to do a hard thing. Then we're also trying to do a hard thing—I mean I think that's all the function of a university—I think we're also trying to do a hard thing by asking these questions.

James W. Houck: At some point, I think it was Claire again, and I'm going to maybe corrupt what you asked and paraphrase it by asking: just because we can do this, do we want to do it? There are a variety of answers that come up I think rebuttals that, one of them being, "well, the other guys are doing it, so we have to keep doing it, because the other guys are doing it." Or, Marc's answer in some form, which is, "well, it gives us convenience, it makes it easier

for us, so we do these things.” Or I think maybe a stereotypical view that I might attribute to an engineer which is, “we do it because we can. Simply because it’s there to conquer.” I’m pretty sure that I don’t have any more answers today after hearing all this than we had going into this.

James W. Houck:

I’m hopeful that this event has made us more thoughtful about this and more curious about all of it. That those of us who are late people, don’t leave this to the experts, because in some ways this is the big issue that’s coming in our time and how we reconcile all these things. I’m really grateful that—again, coming full circle to a thank you that I gave last night to our panelist and to our moderators—really, really grateful that you’ve given your time to this project. For those of you who had to travel a long way to get here, that you did that and were willing to do it. I hope we’re all a little bit better for it.

With that, thank you very much. Thank you for coming to Penn State, and thanks for giving your time for the past twenty-four hours. Safe travels to all of you. Thank you.