

How Artificial Intelligence May Affect the Landscape of Social Security

- Part 1

April 17, 2024

Marcela Escobar-Alava, Deputy Commissioner for Systems in the Office of the Chief

Information Officer Front Office: All right. Welcome. Good afternoon and good morning to those of you on the west coast. Thank you for joining us today. My name is Marcela Escobar-Alava. I am the Chief Information Officer at the Social Security Administration. I have the pleasure of welcoming you to our 26th National Disability Forum, "How Artificial Intelligence May Affect the Landscape of Social Security, Part 1." On behalf of Commissioner O'Malley, SSA executives, and everyone at the Social Security Administration, we hope all of you are well. Let me start with reviewing some housekeeping items. First, I want to inform everyone that the National Disability Forum is a public forum that may include representatives of the press, so any statements or comments made during the forum, may be considered on the record. This virtual forum is being recorded and will be available on the National Disability Forum's website within four weeks after today's forum. Second, we have disabled the chat, microphone, and video feature for our attendees. If you dialed into this Microsoft Teams Meeting, please use your phones mute feature. Third, we are offering two accessibility features today. We have an American Sign Language interpreter and closed captioning. If you would like closed captions please go to your MS Teams toolbar, select the three dots titled "more," select "language and speech" option, and select "turn on live captions" option. Now, it is my honor to welcome and introduce the Commissioner of the Social Security Administration, Martin O'Malley. Martin O'Malley was nominated by President Biden to be Commissioner of the Social Security Administration and following confirmation by the U.S. Senate was sworn into office December

20th, 2023, by Senator Ben Cardin. Prior to joining SSA, he served as Governor of Maryland from 2007 to 2015, following two terms as Mayor of the City of Baltimore. A pioneer of using performance management and customer service technologies in government, Commissioner O'Malley has written extensively about how to govern for better results in the information age by measuring the outputs of government on a real-time basis. Commissioner O'Malley graduated from Catholic University in Washington, D.C., and earned his law degree from the University of Maryland School of Law. Commissioner O'Malley is a life-long public servant. His unmatched enthusiasm for government service and unwavering dedication to public welfare are truly remarkable. Countless individuals, including myself, are grateful for his exemplary leadership. Commissioner, it is my sincere honor to welcome you to our first National Disability Forum. The floor is yours.

Brian Peltier, Chief Architect for the Enterprise Architecture & Innovation Staff (EAIS):

Thank you, Marsella, I just want to jump in.

Marcela: Okay.

Brian: I believe Commissioner O'Malley is running a little bit late.

Marcela: No worries.

Brian: So, I'm going to go ahead and jump in and kind of start to kick us off. Thank you, Marcela. Appreciate it.

Marcela: So, let me just introduce Mr. Brian Peltier who is my Deputy CIO at the Social Security Administration, and he is also actively working on all things AI for the Social Security Administration. Brian, the floor is now yours.

Brian: Thank you, Marcela. As Marcela had mentioned there is no chat feature available for you to make comments today. I just want to make you aware that if you do have a question about

panelists or provide a comment, you can do so via email at nationaldisabilityforum@ssa.gov, all one word, nationaldisabilityforum@ssa.gov. Please include your name and your email to make sure that we know who to attribute the question to. When submitting a question, please do not include any personally identifiable information, such as Social Security numbers or address. We are monitoring the in-box throughout the forum. We'll share questions with the moderators as time allows. If your questions are not answered during the forum, we'll make every effort to answer your questions via email after the forum or share them with the appropriate authors. In a few weeks we will post a link to the recording on our National Disability Forum website at www.ssa.gov/ndf in the outreach section under today's tab, 04-17-2024. As with all National Disability Forums, today gives you, our stakeholders, an opportunity to share your unique insights directly with us and policymakers within the agency. The NDF is not intended to be a means reaching an agreement on an issue and SSA participation is only for the purpose of getting insight by listening to the panelists and their responses to your questions or comments. By sharing your thoughts and experiences you will help us shape the future of Social Security by strengthening our disability policy development and contributing to our continued efforts to address equity within our disability policy and practices. I hope you are just as excited as I am to hear from our expert panel. As a reminder, this is part one of a part two series. Part two will be on Wednesday, May 15th. The purpose of this forum is to learn from our expert panelists, and how you, how artificial intelligence may affect the landscape of Social Security. During today's forum we will focus on what is AI, advancements in AI, challenges, and opportunities of AI on customers and representatives. SSA has been using artificial intelligence for over 20 years to help us create more efficiencies to improve our customer service. With this example we take great pride in making sure we do not impact the rights of our customers. Some examples of how

we use AI today is around our quick disability determinations product which helps us use a computer-based predictive model to screen initial applications to identify cases where our favorable disability determination is highly likely and medical evidence is readily available. By identifying QDD claims early in the process, we can prioritize this workload and expedite the case processing. Another example is in our workload, massive disability workload, is Imagine. Imagine is an AI tool that reads through and summarizes key information within the many, many pages of medical records sent to us for disability claims. This helps staff review and understand medical evidence faster getting to a better outcome for our customers. Critically, both tools are decision-support tools only. They do not direct any particular outcome on a disability claim. That decision remains squarely in the hands of our adjudicators. Rather these tools are simply designed to help them with their work. To get us started for today's forum, the moderator will be Gina Kline. After the presentation there will be an open question and answer session. Anyone is welcome to submit for this panelist via NDF EML box mentioned before, nationaldisabilityforum@ssa.gov. Gina Kline is the founder and managing partner of Enable Ventures. She is an investor, an entrepreneur, a widely recognized civil rights lawyer, and a thought leader. Prior to launching Enable Ventures, Gina founded Smart Job, LLC. This is the world's first impact investing platform dedicated to ecosystem building and scouting the earliest-stage companies in the disability technology market in advancing Evergreen Impact investing opportunities. Mrs. Kline was one of the earliest players in the global disability technology ecosystem to advise, coordinate, and invest in companies through an impact investing approach. She is dedicated to building the future of work by advancing the rights and interests of people with disabilities as innovators, workers, business owners, and consumers. Mrs. Kline is singularly focused on closing the disability wealth gap through impact investment. To learn more

about Mrs. Kline and our expert panelists and SSA executives, please visit the National Disability Forum website. Click on National Disability Forums from the right-side menu, then select the 04/17/2024 tab. Again, that is 04/17/2024 tab. We would like to extend our sincere appreciation to Mrs. Kline and all the panelists for their participation in today's discussion.

Ms. Kline, we welcome you. The floor is yours.

Regina Kline, Founder and Managing Partner of Enable Ventures: Thank you so much, Brian. And I see Commissioner O'Malley has joined us. I might just turn the stage back over to the Commissioner if he would like to say a word or two.

Martin O'Malley, Commissioner of the Social Security Administration: Sure. I'll be super brief on my stage. I mean, the bottom line is this: Our agency is all about serving people and to the extent that we can take advantage of new technologies, new automation, artificial intelligence, in order to better serve people, that's a good thing. We are now in the midst of a customer service crisis. And to this crisis we have been able to attract some really outstanding people. And you just heard from Marcela, our new CIO, and her ability to innovate, her ability to iterate, her background in product development and process improvement, combined with Brian and Mr. Lemmon's, you know, experience in this agency, I think means that the days ahead are going to be not just challenging for SSA, but perhaps transformative. There is already some really forward-leaning things we are doing with Imagine, more things on the horizon. And so all of you are the experts. I'm going to be quiet. I get to be the one generalist in the family. I'm like the three-cord wonder. I know three cords, but everybody in the band, as Marcela would say, knows how to play an instrument. And certainly, we have an outstanding team in the office of CIO, and we are looking forward to learning from all of you and the ideas that you bring forward in this forum today. Thanks a lot.

Regina: Thank you so much, Commissioner O'Malley. And thank you to the agency for such a warm welcome to this event. And I just want to join SSA in welcoming our panelists today and attendees to this National Disability Forum - How Artificial Intelligence Might Affect the Landscape of Social Security. And, of course, this is a part one in a two-part series. It's a very exciting opportunity to discuss these topics. Today, we have a closed discussion with four panelists that are absolutely experts in their field. After the discussion, we plan to accept questions via email. And as time permits, I will share those questions with the panelists. So, we will go out to poll your questions and submit them to the panelists as time allows. If you wish to ask a question or provide a comment by email, please include your name and location in the email. The appropriate email address for today's questions are nationaldisabilityforum, all one word, @ssa.gov. That's nationaldisabilityforum all one word, @ssa.gov. And the chat line will not be open during the discussion segment so all questions must be sent via email to that address. Before we begin this afternoon's discussion, it's important to notice some of the themes of today's event. You heard a press towards improving customer service, reducing delay, or creating transparency and access to government programs. There are some parts of this discussion around technology that endemic and inherent in conversations of technology have to do with avoidance of discrimination, avoidance of introducing bias, protecting personal information, uplifting the consumer experience while safeguarding security. So, we are excited to get into these conversations with some esteemed panelists and very excited to transition now into the discussion of how artificial intelligence may affect the landscape of Social Security. And let me introduce our panel: Rylin Rodgers is the Disability Policy Advisor with Microsoft; Adler Archer is an Associate Research Scientist in Biomedical Engineering, Director of the Office of Strategy Management, and Managing Director of the Inclusive Innovation Initiative at Johns

Hopkins University; Daniel Schwartz is a Fredrickson and Byron Professor of Law at the University of Minnesota Law School; and Jason Green-Lowe is the Executive Director with the Center for AI Policy. There are rather extensive and impressive, I might say, bios of these executives and panelists on the National Disability Forum website which can be found at ssa.gov/ndf. And you might want to check that out. Let's begin. Let me start with a question to the panelists and maybe I could start with Rylin if you'll let me. Might you provide us with a simplified explanation of what AI is for those who may not be familiar?

Rylin Rodgers, Disability Policy Advisor with Microsoft: Yeah, thanks so much. I think that's a great place to start. We all spend a lot of time thinking and hearing about AI these days, that term artificial intelligence, and I think it's really helpful to have a long-term and big-picture perspective of it. So, artificial intelligence isn't something new. We have been talking a lot about Generative AI, which is the latest iteration of artificial intelligence, but think it's important to understand that artificial intelligence dates back to 1956 when the field of computer science was created to think about intelligent machines that could replicate or exceed human intelligence. Then we fast forwarded to 1997 and we got to that concept of machine learning, that subset of AI that enables machines to learn from existing data and improve upon that data to make predictions. In 2017 we really think about deep learning which is machine-learning techniques in layers and networks to make that process of data decision-making much faster. And then we are all in this moment of generating AI, where there is the ability to have plain language, written, and visual, and auditory content through large language models. I think all of that is important particularly when you think about disability and accessibility and the ways AI can be used because it's really a combination of different types of artificial intelligence that we're benefiting

from and that make our lives and potentially make processes in government systems work better. So, it's not one thing or the other, but often the totality of AI.

Regina: That's really interesting and very well said. Let me turn it over to Jason for a moment. Jason, I'd like to ask you what are the main types or categories of AI and how do they function differently?

Jason Green-Lowe, Executive Director with the Center for AI Policy: I think Rylin gave a great summary of the technical categories of AI in terms of how they affect people in practice. Over the last couple of decades, we have had classifiers that mostly sort an object into one of a few different buckets. A great example is a spam filter. All it has to do is decide for any given email is this email spam or not spam. So, they are not going to create new content. As Rylin said, today we're in the moment of the Generative AI, we have AI that is creating new material, whether that's poetry, or sums, or images. And over the next five years, I think we are likely to see a transition into Agentic AI. AI that can independently pursue a goal and make something happen in the real world. So instead of just writing a travel guide, an AI agent will be able to connect with robots or with mobile apps and actually get the items on your to-do list done, go buy a flight, buy a hotel, buy bathing suits for the dates of your trip, sort of make it all happen just in response to a prompt, the same way today a Generative AI could create an image of response to a prompt.

Regina: And, Adler, following up on those comments from Jason, can you give examples of everyday applications from hospitals to insurance companies, of AI that people might encounter without actually realizing it at all?

Adler Archer, JD, MS, Associate Research Scientist, Biomedical Engineering, Director, Office of Strategy Management, Managing Director, Inclusive Innovation Initiative, Johns

Hopkins University: And, so, one really popular area is around diagnostics. So, radiologists use AI software to analyze images and it has been proven in a lot of cases to be just as good or sometimes better than humans at detecting tumors and other types of anomalies and healthcare. And there are lots of other ways that we kind of think through this. One big consideration is when software is acting as a medical device and so the FDA has a lot of regulations around if something is treating, diagnosing, preventing illness. And so certainly there is a lot of consideration there as we start to look at new ways to deploy AI. Virtual assistants are becoming more common where people can kind of ask pretty basic questions, again, that aren't crossing the line into actually diagnosing, but instead maybe providing medical information that can be helpful. And then robotic surgery is another one that comes to mind and so we have a technology that came out of one of the labs at our school that's a robotic device that can assist surgeons in performing procedures remotely so that they are not necessarily even in the room where it's happening.

Regina: And, Daniel, are there any misconceptions about AI that you often encounter, and if so, how would you clarify those misconceptions?

Daniel Schwarcz, Fredrikson & Byron Professor of Law Distinguished University Teaching Professor, University of Minnesota Law: Well, I think one really is this idea that AI is this all-encompassing area and so a lot of the risks that we think about for AI depend upon both what type of AI we're talking about and how it's being used. And so, I think there is this tendency to go, well, all AI is biased, or all AI is sort of this machine that is acting completely outside of the scope of human control. And the truth of the matter is that there are a lot of nuances in how AI is used; what AI means; and what types of risks it would pose. And so just to like be a little bit specific here, right, one common use of AI that occurs in the insurance field though I don't think

was mentioned in the Social Security field is fraud detection. Now, fraud detection uses to go sort of the basic categories, it usually uses machine-learning AI so there is no one that sort of comes in and says these types of claims we should flag for fraud. Uses machine-learning eyewear, that machine is sort of fed examples of past fraud to predict future fraud. Now, that's something you might imagine maybe in the future Social Security Administration would be interested in doing, but there are also a ton of risks that come along with that. And to illustrate that point you can have biases in past cases where you found fraud, where you particularly are likely to find fraud among certain populations, but not other populations. The AI cannot only replicate that but supercharge that. So, in the specific context of fraud detection or you are using machine-learning AI that is predictive in order to tag people for fraud, there is a risk even if there is a human in the loop because once you are investigating someone for fraud, there are already real consequences even if ultimately at the end of the day you decide there is no fraud. So, I think one needs to be really careful in, if I am going to say three or four things to specify, what the AI is that's at issue; how it's being used; what the risks are that are in play; and how those risks are being litigated.

Regina: And so, this is a natural consequence of your response is for me to ask Jason what does regulation play in all of this?

Jason: Yeah, so I think a lot of the regulation so far is tied into Daniel's ideas. It's use case specific. So, you have the FAA regulating the use of AI in airplanes; the FDA regulating the use of AI in pharmaceuticals, and so on. The FTC is starting to do some general-purpose regulation by dint of their role in consumer protection because at one time or another all of us are consumers. But there is no Office of AI, at least not as a regulator, that deals with general purpose AI threats that cut across industries. And while it's true that AI isn't a monolith, not

every piece of software behaves exactly the same, there are some commonalities, at least as the Center for AI Policy sees it, in terms of generally AI being able to expand in ways that could be dangerous, that could pose unanticipated cyber security risks or unexpected proliferation risks. And many of those risks are going unregulated because they don't fall into an existing use case that neatly fits under a particular industry regulator's umbrella.

Regina: And on the flip side of that, you know, that you just addressed, Jason, the current considerations around regulation and future considerations around regulations. And Rylin, how could we talk about the current need for responsible and ethical deployment of AI while minimizing potential risks and biases, meaning, how do we capture the opportunity set while minimizing risk and potential introduction of bias?

Rylin: There's a lot there. A couple of things I would raise. One piece on the regulatory side. I think it's helpful when we think about responsible and ethical AI to understand that the impacts, outcomes, and use still are responsible to existing regulation and law. So, when we think about potential concerns related to bias and discrimination from AI, it's important to know that users are responsible, too, for existing civil rights law, whether it's the ADA or other pieces. Just because AI is involved it doesn't mean that there are existing rules of our society are no longer applicable. So, I think that's a really important piece. At the core, responsible AI is incredibly important and it's been helpful for us to ground our efforts at development in four core pillars really thinking about fairness, reliability, and safety, privacy, and security, and inclusiveness, you know, that ability to make sure that when we are thinking about AI, we are thinking to ensure that it treats everyone fairly and avoids affecting similarly situated people in different ways. How do we think about that bias piece. And exclusiveness is particularly important when we think about disability issues. AI, when we think about large language models, is trained over

the world as it exists. And the world as it exists includes significant bias. We see racism, sexism, and ableism. But there are ways in developing your model and deploying products to really protect against reinforcing or accelerating those biases. And it takes attention to that. It takes attention to that red teaming, to setting parameters, and it will take really careful attention to what are the datasets being used, what is underrepresented, how do we ensure that it reflects who we really all are. Reliability and safety is critical for AI to be effective and for us to be able to use it in government and other settings. We have to trust it. We have to know that when it's operating safely so those stories and concerns about, you know, when AI starts working outside of its needs. So, what are those safety regs that could be part of deployment? I think those principles are starting to guide practice across the corporate sector and then we are seeing it reflected in some of the first attempts at regulatory efforts, the executive order, things happening in the EU, et cetera. So, the conversation is rich and continuing.

Regina: That's a great, very thorough response, and very much appreciated. And very much appreciated your point about the current status of U.S. law still applying to technology. Things like the Americans with Disabilities Act, regulations that were, and guidance created by the U.S. Justice Department and the EEOC around bias and discrimination. So, point well taken. So, shifting into advancements in AI, Jason, what recent breakthroughs in AI tech have had the most significant impact across various industries?

Jason: One of the most surprising discoveries about AI is that the particular breakthroughs tend to matter less than the sheer volume of resources to get it thrown at the problem. So, there is new exciting research projects being published all the time, but in the time, it takes you to do that research project, the amount of data and the amount of computing power that you can buy to implement the results of your research could double or even triple. And those extra resources

usually make more of a difference than whatever you learn from your research project. That's why we encourage people to think about progress in AI in terms of a curve. There is an exponentially increasing curve where every year AI is a little bit smarter than the year before and the rate at which that improvement is happening is speeding up, and I think one of the results that we're seeing of that is that companies are increasingly rushing to adopt AI within their operations not just to compliment the people who are working in their roles, but increasing their planning on what's going to happen in the near future when AI starts substituting for some of the human labor.

Regina: And, Adler, how can we ensure that -- I think Rylin sort of set the stage for you to answer this almost perfectly, but how can we ensure that AI development is indeed inclusive for everyone. That it exercises a lens that includes from the beginning the manifold diversities of our community and our society?

Adler: Yeah, that's a great question. So, I think about three different things: So, the first is data to algorithm because in a lot of ways the conversation can focus on what the technology needs to do, but it's really considering the fact that the technology is getting the data from humans who have bias and so there needs to be some type of attention to the humans who are putting the data into the machine in the first instance, and then the machine is going to do whatever it does which may or may not add bias which would add another layer potentially of bias on top of that. But then that's coming back oftentimes to a human who is going to use that product from the AI or the machine that's created it which could add a whole other layer of bias depending on the human that's receiving it. So, I think that technological intervention is important, but also there is the associate technical doctor to consider around how people perceive the information they are receiving and then what they do with that in terms of how they engage with the computer. And

so, you can see that in terms of representation bias when you are thinking about data that we might put into a machine, are we considering a diverse sample of people or are we focusing more on a certain cohort of people because that's going to get different information. And then, you know, who's analyzing that data on the output because they may have different experiences, or different expectations, or different ways that they see the world that are going to impact what they do with the data. So, in terms of really having a fair outcome, I think you have got to consider all of that.

Regina: And would that lend towards a world view that would support more diversity in tech and in particular the people that build tech, and not to show all of my biases here, but how would you respond to that comment?

Adler: Yeah, definitely, I think as far as diversity and tech it's important when you think about diversity and leadership roles as well. So, we have a grant from NIH around developing neuro-technology and kind of what we explained to them when we submitted our proposal is that we have three levels that we think about it. So, part of it is who is doing the innovation. Certainly, one diversity includes a group Gen innovation, but who is represented as an expert in this space because sometimes people have a really limited view of what an expert looks like or who an expert is. And so, it's important from a representation standpoint to have a really broad set of people who are represented as experts. And then I think the third, and probably most important to me is who are we trying to serve with the technology. And I always say, you know, Hopkins, we're really good at making shiny expensive things but there are lots of problems that don't need a shiny, expensive thing as a solution. And so, who is thinking about solutions that aren't necessarily going to generate a lot of revenue and so I think that's important to who are you trying to serve.

Regina: And, so, Daniel, what do you claim adjudicators need to know about how AI summarizes medical records? It is certainly very pertinent to the conversation around Social Security.

Daniel: Yeah. So, when we are talking about summarizing medical records, we're often talking about these who generate AI, and in particular, largely responsible. So, this is one of the most, I think, you know, widely distributed use cases for professional settings, for LLN's. And so essentially what you can do right now is you can take a lengthy document and you can either plug it directly into something like chatGPT, or one of the other large language models, or you can use a system that includes various safeguards from a privacy perspective in order to generate a summary. And there are technologies that are sometimes called retrieval augmented generation that allow for this to happen in a relatively automated fashion where you might ask a question about a document, or a series of documents and the system will retrieve those documents and then essentially use the AI to search them and summarize key points. You know, I think that there are a few key things to know. So, the first is that there is very good evidence that when you are dealing with retrieval augmented generation when you are providing the underlying source material that is being summarized, LLMs are very good. They don't hallucinate nearly as much hallucination as this term for making things up in a very confident fashion, something people accuse lawyers of all the time. And it turns out that sometimes AI can do it as well in generative AI. But it does it much, much less when you are dealing with retrieval augmented generation or summarizing with the medical records or other documents. On the other hand, it can happen. And so, the tricky thing is how do you manage a system that is -- can work very well already, perhaps even better than many humans, but at the same time certainly is still not perfect. And I think that the best answers we have right now are being very aware of holding those realities in

your head, but this can be an incredibly useful tool that at the end of the day we have to recognize humans make mistakes, too. And if you are going to ask a human to summarize a medical record, they very well may make more mistakes than an AI. So, what that means is that I think we just constantly need to be empirically testing this; right? We need to have data. We can't make assumptions because AI's, you know, the effectiveness of summarization technique is going to depend on how it is implemented, how it's prompted, who the comparison group is in terms of the human who is summarizing. What is done with that summary? Is it then checked? So, there are a lot of issues there and I don't think that one can sort of simply say it either is a good idea or a bad idea. But I guess what I would say is it is absolutely clear to me that this is already an important tool in the toolbox of many professionals, and I would include anyone who is frankly adjudicating a claim for Social Security or any other, you know, insurance like claim to have this as a tool. But using the tool correctly, I think we have a lot more work to do before we figure that out. And the difficulty is it is going to be always changing because as mentioned earlier, the AI is already changing radically, you know, on a short-term basis and in an exponential fashion.

Regina: So, that was mostly a conversation about summarizing medical records. So, if we stepped it up to the next level, the next question, we would say, wait a minute, after this evidence is reviewed can AI responsibly assist in making determinations? And I would ask that question of Jason. Taking what we know to be true about AI in reading large sets of data and information, and what we just heard about a hypothesis about whether it can evaluate medical records, can it actually make determinations at some point?

Jason: Technically, I think that point is very rapidly approaching; right? So, you already have the ability to visually decode the blurry images, make sense of ambiguous written text, mimic

human responses. In some ways evaluating a case file within a formal bureaucratic process is an ideal scenario for an AI because it's going to get repetitive structured inputs and the AI will have an easier time kind of classifying them drawing on some of the tools of machine learning that were developed over the last few decades. So, the main challenge is ethical rather than technical. It's not so much a question of can AI make thoughtful recommendations about how to rule on a case file but should it.

Regina: Well, let me actually turn first to Adler and then Rylin on that issue. You just raised a question: Should it? What are the ethical considerations, if AI is helping or supporting a determination process, what ethical considerations should we go to first? Adler first and then I will go to Rylin.

Adler: It's funny, because we were talking about some of the myths that people have about AI's and as a biomedical and informaticist to have this conversation with doctors a lot, there is a concern that AI is coming for their jobs and it's going to be able to do everything. And so, I think AI is best suited as a copilot, and there's are lots of things that it will do well, but we kind of always want to have that human pilot there to make sure that things are operating smoothly. So, that said, I think cultural competency is an important area to consider, so thinking about who is being engaged. Are there language barriers, are there other things that are going to impact the way the person receives the information to making sure that it's done in a way that is customized to the recipient of the information. And then I think transparency. So being really clear when AI is being used so that people know that they are interacting with AI and not a human. I deal with that when I call the banks sometimes, I go, I don't want to talk to the computer; I want to talk to a person. So even just this kind of thing being really clear with people so that they have a sense of what's happening. So, I think those are some basic things that are a good start.

Rylin: Yeah, this is Rylin. I would build on that. We agree that AI is a copilot, and in fact, that's how we brand our AI tools. And I think that's part of the ethical consideration. It's not whether AI should be used, but how AI should be used to support the work of reviewing medical records and helping it to get to decisions. We would, you know, say, that it's not the answer to do those pieces, but instead a tool that can be moved forward. And we're learning about how it can be used effectively. One thing that is incredibly helpful in the summarization is you have the ability within tools to get a citation, where is this information coming from in the original document. And then that gives the human reviewer the opportunity to check, to confirm. And also, a summary has the ability for a human to ask questions of what when something is missing, when there is other information needed, you can have that back-and-forth engagement with materials using AI to assist in that way. And then similar to the making the decisions, AI likely should not be making a determination decision, but can help generate what is the evidence in the file that supports this decision. How does that crosswalk to what the law requires, what the categories are, and give that first draft of documentation to allow the system to go back and see are there critical pieces missing? Was the decision based on information that is not fully reflective? That piece of the heavy lift to get to the draft document to allow the real human process to move forward in the determination process as regulations require. Is this significant, accelerated or in time, which gets to something that we all desperately want which is a system that works better for the individuals who need to be served. I would also add there is opportunities for it to serve individuals who are often overwhelmed by the way systems work. Wouldn't it be lovely if an individual can get a copy of that summary and had a chance to in plain language look at what it says about their medical record, then they can look and see does that match? Is there something missing? Let me, oh, this is first for me, then I need to get this additional piece. And the other really exciting piece

in terms of that is there is the ability to take that information and give language access in multiple language which is incredibly important to reflect who needs to engage with systems in America. The other piece that I would echo, and I love that point about transparency, when is AI being used? Another important piece of the epochs of it is what are the datasets being used and we should all be honest about that. When we're thinking about claims adjudication, it should be very easy for the public, for any individual involved in the system, to know what the AI system is using to look through and help make those first drafts, or whatever the process is. So, I think that transparency, not just the "when," but the "what" is incredibly important to the ethics.

Regina: Let me just transition into a follow-up question staying with you, Rylin, just for a second, it seems like there are both challenges and opportunities from a customer side and a representative side in the Social Security system, leveraging AI. In terms of what we heard from the Commissioner, we were talking about delivering, improving customer service. How can AI be leveraged to enhance a customer experience that actually maintains a human touch, meaning that we're not losing a grasp on who these systems have been set up to serve and how these systems have remained or should remain more accessible to the end-users for whom they have been built since the new deal. So, what is it about AI that could be leveraged in a very modern way to carry on a very old thought which is that these systems were created to serve customers who are eligible. What are your thoughts about that?

Rylin: I think that's exciting to think about. Through the realities when we call and want to speak to a human, it's often because the other information that's available to us is not easy to navigate. I don't think anybody has ever had a delightful time reading through the regulatory process related to Social Security. If you've had a delightful time in that I would like to meet you. So, people want to talk to a human that can help them understand that. But AI gives us an opportunity to

take the information that somebody needs to understand what's available to them, what they might be eligible for, and how to move forward and create a plain language easy-read summary that gives them a place to start, and that's a huge step forward. The other piece that I mentioned before that's delightful about that is that you could instantly translate that into multiple languages. I also love the ability to give people a first round of asking questions through at the AI tool. That doesn't mean that they might not still need or prefer to speak to a human. But many of us would feel more prepared for those detailed conversations if we were able to get the information, we needed in a format that made sense to us. So, I think that's a really important piece. The other part about this that we haven't raised yet equity issue here in the United States. And we're talking about Systems and AI and technology and digital access, and there are still many citizens for whom digital access to systems is not part of their daily life. So, when we think about those tools moving forward and how they can support sort of customers in the general public and also systems, we need to acknowledge that that is not currently reaching everyone in America, and how do we account for that in terms of that human touch of customer service in a system like SSI.

Regina: That's a terrific response and I appreciate how thorough you were in addressing the human touch. So, Daniel, what are the main challenges faced then by customer service reps in, you know, adapting to AI power tools? If I'm a customer service representative trying to assist a customer, what might be some of the challenges?

Daniel: Well, I think that one of the real challenges, frankly, is that there is evidence that that is the sector where there are many use cases where automated tools can operate as well as many customer service agents. So, there's some empirical evidence demonstrating that at least when it comes to newer claims, you know, customer service agents, completely automated systems, or

nearly completely automated systems, can work better. Now I think that in most cases what we're seeing is either that, like humans we are using the AI to help them generate solutions. But there's a question about whether or not at the end of the day we're going to have a first-layer system, I think we already do to some extent, where you're interacting with the AI and then you can bump up to a human, and I think there will be changes where the automated system works better and better. So, right now, I think that the question is, how do you use these tools effectively, but I also think that there are real questions about how we should change our customer service models, because in some ways, those sorts of customer service experiences, I think even, a more idealized setting for generative AI than many others. It's one where we have a lot of really good empirical evidence. So, I think that we need to embrace some flexibility and change in that setting.

Regina: And Adler, I would like your opinion, how do you think AI-driven customer interactions can be empathetic, personal, or not lose the human touch, or do you just think that they can't? I mean, where do you come out of this?

Adler: Yeah, I definitely think they can. So, part of it is when we're talking about this idea, AI as a co-pilot earlier, so are there other ways where even in the transition from AI to human that it's clear to the human you're transitioning to what you've done with the AI, so the bank analogy I gave you, I'm pressing zero like 15 times trying to get to a person. So even the person being aware of like my interactions with the computer before I got to them could be really helpful to make me feel like the investment with the machine wasn't a waste of time. There is some type of connection there. And, you know, emotional intelligence, which is something even as humans, you know, we have to really work on developing. I think the same is true with AI. Having an empathetic response, if you're dealing with something that's really personal instead of something

that seems robotic, even just the tone of the voice that's used, the way that options are framed, I do think there are lots of ways that AI can be more empathetic. The other thing is the ability to kind of personalize because it can take such big datasets and drop conclusions. I think it's really powerful versus my little brain trying to like, really come up with very personalized responses to someone having access to, you know, thousands and thousands of claims and looking for trends might be able to really, actually better empathize with what the person is going through. So, yeah, I think it's possible. It takes time and that has to be the intention. The intention has to be an emotionally intelligent response, not just a product of the AI tool.

Regina: And as you said that can often be a challenge in even humans in customer service roles but noted on both levels the importance of it. I want to ask Rylin what opportunities does AI present for improving customer engagement and satisfaction across various channels; and by channels, I think I mean Chatbot to Virtual Assistant? So, interoperability across a platform, how do you respond to, you know, the opportunity for really engaging customers across multiple strategies of channels of communication?

Rylin: Yeah, I think there's a lot of opportunity there, and one of the things that excites me is about the opportunity to enhance accessibility across the dimensions of disability, and we're learning a lot from that firsthand. Microsoft has something called a Disability Answer Desk. I think it's one of the best kept secrets in terms of customer support that exists, but if you're an individual with disability accessibility need and interacting with Microsoft technology and running into a challenge and need help, our Disability Answer Desk is here to help. But it gives us a natural place to learn from what is a -- what are the AI opportunities and how can they best support our customers, and individuals around the world. And a couple of things, we do have a chat feature where you can chat with a human. We also have a Chatbot feature, or an AI feature

which allows you to ask those questions. And we see people do both sometimes both simultaneously, sometimes getting sort of a first round coming to a second round. One of the things that's incredibly exciting in terms of enhanced accessibility, we have an option to engage with us through a partner tool called Be My Eyes and Be My AI. And this supports access to the information and support for individuals who are low-vision and blind allowing them to get high-quality image descriptions as to what their technical issues are to work directly with our technologists to resolve them. And it really had rapidly accelerated how you can share information and I can see something like that being incredibly helpful in terms of the documentation and paperwork that people are navigating if you have hard copy, old-fashioned medical records, and you're a blind or low-vision individual how do you navigate what needs to come to SSI, those kind of pieces and parts. You know, obviously, thinking back to other parts of AI and captions and the ability to have both ASL and caption video calls and moving those pieces forward is really important. So there's this way to really think across the [inaudible] of customer service options and I think that's an important piece, because it's not that AI is now the only answer in resolving all issues, but it's an important piece and really critically, it's a piece that can allow to unblock accessibility issues and accelerate the needs of people with disabilities getting met and really cutting down that timeline and getting to that equity piece. So, I think those parts are helpful for us to think about when we're thinking about customer service and the option. I think the other pieces referenced here today is this is what is available right now; what's available tomorrow, like literally tomorrow, may be very different and really being open to being those early adapters, which disabled people often are, in terms of utilizing it to meet a need and really thinking that through in the customer service space.

Regina: Yeah, that's really noted and terrific response. Jason, there's a tension between full automation and human intervention. And we so often, it's so en vogue to have conversations in technology circles around Human in the Loop. Human in the Loop as a safeguard in Quality Assurance over tracking AI. How do businesses strike the right balance here, and by the way, the right balance as applied to customer service? So where is there a role for human intervention and quality assurance and automation and where do you come out on balancing those two things?

Jason: Yeah, one of the most important things that a Human in the Loop can do in customer service is gather information about where in AI is not meeting the needs of customers. So they are almost more fulfilling the role of a customer service manager than a frontline customer service representative because an AI isn't necessarily able to explain its decisions or to explain how and why it communicated in a particular way, so you're going to need someone who can ask the customer, hey, what went wrong here and why, because the AI isn't necessarily in a position to explain that to you. You can collect data about different kinds of encounters, but even the AI may not understand why it said what it said or be able to give an accurate accounting of that. In terms of figuring out whether you've reached the right balance as an agency or as a business, I think it's important to ask the right questions and to have a truly open-minded decision-maker. So, when you evaluate the success of a partially automated customer service program it's not enough to just ask, okay, did this automated service solve your problems or to see how many people were able to use the system through to completion, you really want as Daniel was saying, some kind of data on who's achieving their goals. So not just who completed the process or who stopped calling to complain, but who achieved the thing that they set out to do in the first place with the aid of the AI, how common is that with and without the AI assistance? And the person looking at that data and deciding if it's adequate to support the decision to continue the

automation, probably shouldn't be the same technology champion that was advocating to install the system in the first place. It's only natural for the people who are most excited about the technology to take a lead role in designing and implementing it, but then you need to make sure that there's someone technically literate in the organization who is not that initial champion who can look at it with a fresh pair of eyes.

Regina: And Adler, what strategies can organizations implement to both train and up skill their customer service representatives to work alongside the AI? So, what is the set of skills necessary both in understanding the core technology but understanding principally how to work with that core technology in customer service?

Adler: I think two things come to mind for me and I think about this in hoc literacy as well, so in this instance, information literacy, technology, terms, how these things all work, because we see in you know, popular media kind of what AI is, but it's sometimes hard for people to really understand exactly what we mean, and even how that's different from robotics. So, you know for some people to think AI, they think the terminator, and so, maybe that's hopefully much further down the road. So just helping people get on the same page about the terms I think is an important part. And then also, helping people understand what their role is in all of this. That's such an important part of training, right, is for people to be clear on the distinction between what they are doing and then how these tools can help them do their jobs, because at the end of the day AI is just another tool. So, I think taking that approach [inaudible] a competency that they can develop at different skill levels, and then helping people see how it's important and helps them do their job better or faster or in some other way that there's a value for them. I always think about my friend, he was in the Marines, always got all of these acronyms, but he talks about "what's in it for me?" So, I think the same is true, I would say for the representative, and

then also what's in it for the customer because a lot of people would be motivated by knowing these tools can help them provide better care to their customer.

Regina: Well, and you're talking about being motivated by what's in it for me and you can imagine particularly in very large organizations with a seminal number of claimants and an enormous backlog of cases that streamlining and cutting inefficiency is a central motivation for using technology. I ask Rylin this actually about streamlining internal processes and empowering customers by streamlining, what benefits could there be with the power of this new technology that we all are especially lucky to be alive now where there is technology that can help streamline processes?

Rylin: I think that the exciting piece there is we have this ability to take huge amounts of data and really review it in a meaningful way where previously we were aware of the backlogs, we were aware of the piles and piles and reams of information to get to, and now there's a chance to really take that in and then start to evaluate and move forward in a meaningful way to know, pretty quickly, you know, of the backlog. What are duplicates, how many things are multiple persons same file because it was started in a different place or relocated during the pandemic, or a different disability was acquired? Those are all realities of the overwhelm system. So, I think being able to use the tools to really get our collective hands around the data and then to get it to a place that we can move forward. I think it also gives us a place to think about our own structures of what is the decision-making. How do we now have a greater insight into what the they're/their is and how do we know how to make decisions about that. And it is absolutely about the backlog, but it's also that we have the opportunity to use the data on hold to anticipate what's coming next, to think about and plan for what geographies in the country, down to frankly the neighborhoods where there will be a higher need or different types of support. When we think about those equity

issues that we have referenced, what communities need professionals with different linguistic backgrounds? What communities will not be able to engage digitally because there continues to be a gap in connectivity? And really using that holistic data to make decisions that allow an agency to pivot and function work actively.

Regina: And Daniel, what does really transparent very accountable, incredibly aboveboard AIN customer service look like?

Daniel: Well, I think the most important principle is you need to know if you're interacting with an AI or not, and that may seem like a trivial point because these days it's very clear if you're interacting with a human or not. But actually, you know, we can anticipate very soon that it will be harder and harder to tell. And so, I think customers absolutely need to know if they are dealing with a human or not. I think that customers also need to know if there is a significant reliance on AI by humans, and that is because there are real risks that humans relying on AI will fall asleep at the wheel and that you know, that has been demonstrated in empirical work and it's easy to understand why. Going back to an earlier conversation that we were having regarding claims, and can an AI write a claim out? Technically, absolutely can. It's easy to show that. Now you could say, well, we can have a human in the loop, look at that and decide whether it's right, but anyone who writes, adjudicates claims and who does a lot of [inaudible] will realize that part of the process of thinking through an issue is oftentimes writing about it. So, there can be a tendency to just accept uncritically what you get from AI. So that means that I think there should be a right not only to know that you're interacting with an AI, but also a right to have at least some transparency or capacity to know when a human is themselves relying on AI to make a judgment that is impactful about it. I think the third element of radical transparency is auditing. And the auditing afterwards has to involve not just -- you know again, different uses of AI raise

different risks and require different safeguards. But, in a lot of settings, I think it's really important to audit the AI usage for disparate impacts and then to evaluate those really critically. At the end of the day, you can put in a lot of safeguards. You can try to use good data. You can try to use a predictive model if we're talking about non-generative AI but say AI to detect fraud. But at the end of the day, it's hard to know whether or not there are going to be problems. There are all sorts of stories about people with the best of intentions using AI in ways that create really bad obviously discriminatory outcomes. And so, I think the only way you can really prevent that is auditing, checking what the results are after the fact, not just in terms of satisfaction, but also in terms of potential hidden risks that you wouldn't see. So those are sort of the layers of transparency I would focus on, but I think you need to tailor it to the use case, tailor it to the AI setting. So those are served at a relatively high level.

Regina: I want to ask Jason to look into the future. What are some of the exciting developments and trends that are coming with AI? Where can we expect to be in the coming years?

Jason: I think we're increasingly going to see more or less seamless interoperability. So we spent 20, 25, 30 years building electronic health care records and no two companies agree on what fields of data should be stored or how the databases should be structured. And those questions may stop mattering because we may have AIs that can talk to each other from one health care information system to another, and just sort of produce these summaries that say, hey, here's what I want the next person looking at this medical file to pay attention to, or here's what the next Social Service worker looking at this file should pay attention to from whatever format the data is coming in. And that's going to open up new opportunities for potentially much faster service for the customer base.

Regina: So, I am going to do my homework and I've been told I need to remind the audience of the email address to send us questions, and so this would be as good a time as any - DisabilityForum@ssa.gov. So National Disability Forum, excuse me, NationalDisabilityForum@ssa.gov where we are going to now turn to questions that are in that inbox. So again, NationalDisabilityForum@ssa.gov. And I was going to ask, if the panel would let me, I would love to ask all of you to comment on at least the first question from the audience, and then we can take them as they come in. Has AI been approved for administrative work before, and how well would an AI Chatbot do to screen incoming messages for a complex federal agency? So, I will let anyone that would like to participate first start. Go right ahead.

Daniel: Well, I can say one or two things on that. So, right now certainly there are agencies that are using AI, we even heard that earlier on, to assist humans. And in general, regulatory agencies have pretty broad authority to achieve their statutory aims. So there's no affirmative prohibition in most incidents that I'm aware of, of an agency using AI. And remember, you know, really until pretty recently, AI wasn't what live people were thinking about in setting up say, the Social Security Administration. They were given tasks to perform, they were given authorities, and the absence of any legislation or clear rules on that essentially went to the discretion of the Agency about when, if, and how they were going to use AI without violating any rules. You heard earlier of course that general rules apply. But there's no clear way in which just using AI would violate any rule. So, we did just in the last few months, get an executive order from President Biden and this actual form may be a response for that which was for all federal agencies to evaluate both potential benefits and harms associated with the use of AI and how they are using AI. And so, right now, I think is moment of transformation and change, but I would say that currently there are not a lot of rules. Currently, different agencies are using AI in different ways. I do not think

many agencies are currently using generative AI, to my knowledge, at least at the federal level, I'm not aware of that. I could be wrong on that but certainly at least predictive AI and certain basic uses of AI that have been off ground for a while. There's nothing in federal law that would limit that as long as you are applying the general principles.

Regina: Does anyone else have anything to add to that, otherwise we can move on to the next question. Go ahead, Jason.

Jason: Yeah, so Daniel is right that there's nothing in the way of prohibitions. There is a requirement from Executive Order 13960 that you make an inventory of your use cases. So, if you're going to start using AI as a receptionist or as a routing agent, then that would be something that by the end of the year you would want to include it in that inventory. And there's also some discussion in Congress, should procurement standards for AI and the federal government include some kind of requirement to ensure that the AI being used is safe and secure and trustworthy and verified and so on. And I think we can get there. When I was interning for Senator Bill Nelson in 2004, they stuck me as the receptionist. And one of the things they drilled into my head was if you have any doubt about whether this message should go forward, go find the legislative assistant and refer the case to them. Don't just hang up on someone if there's any doubt at all. And I think we'd want to train AIs receiving messages for the federal agency in the same way. You have to make sure to err on the side of consulting someone to see if this message really does need immediate attention, but it could be done.

Regina: So, another very interesting question from the audience is how are other companies comparable to government agencies that are using AI? So, are there private market comparables in terms of scale or composition that we can look to and see success in applying AI?

Rylin: I think I can start. Obviously, those of us that find ourselves in the tech industry are actively using AI and various dimensions of their work, and we're seeing that -- [no audio] -- so it is, you know, an interesting dimension of the way government and private sectors and all of our experience interact that individuals may be already engaging in systems in their health care space that are using AI to drive pieces that then will have a direct connection to SSI, even if SSI themselves is not using AI. So, the reality of what's currently in place is rapidly accelerating in terms of what's happening across the private sector. I also think it's always valuable to look globally at the -- that the rest of the world is moving and changing as well, and so there are models in other countries, government benefit systems that we can learn from them in terms of how they are using AI as well. So there are sort of multiple pieces moving simultaneously and what is happening in one industry, like technology may be farther -- [no audio] -- but even with an industry that you may presume or have bias that AI is not currently being used. It's increasingly my experience that some sector of that industry is also using AI. So, really understanding how dramatically the landscape is changing and that that's interacting with government services even if the government yet isn't using it fully.

Regina: Yeah, go ahead, Daniel.

Daniel: So, for me there are two industries to highlight that I think are particularly relevant for the Social Security Administration. So, the first is the insurance industry, right. At the end of the day, there are a lot of parallels between what private insurers do in the Social Security industry in terms of adjudicating claims, explaining benefits, having some customer service basis, making payments. And in the insurance base, obviously there's a lot of change. Insurers as I have mentioned several times already, they have for a long time been using AI to help them with fraud identification, though I do not believe that is happening at Social Security. I think that is an open

question, whether or not it should be used, though -- I actually don't know for sure if it's not being used, I know it wasn't mentioned. Where it's also being used in the insurance industry, though here it's sort of a little bit less clear is in the underwriting and rating process, and that's less relevant for Social Security, because it's a government benefit, so that wouldn't be applicable. But I think a third area is in explaining benefits, right. And I think that just in terms of summarizing benefits and communicating with customers, it's being used a lot in the insurance industry, and I think probably less in the Social Security space. And then of course actually pending claims determinations that are specific that form the basis for an appeal and for subsequent adjudication. There I actually don't know how much it's being used in the insurance industry. I suspect it's being used quite a bit, but I don't know, and it's a real open question. So those are sort of four areas. The other one I'll quickly mention is just the legal services industry. I'm actually a law professor and a lot of my work focuses on how the use of AI can help lawyers. And at the end of the day, we've already shown that like some of the older generation AIs can really improve sufficiency in terms of some basic legal work. And so, to me that points to the fact that while I'm not at all sure that we want to have an AI helping make claims determinations, I'm pretty sure that we're at the point where when it comes to writing up an explanation of a benefit denial or a claim, that we should be using AI to help assist with that, that that will improve efficiencies. And so, I think we can already say that although there are of course, you know, ethical considerations and guidelines. I think the evidence from the legal sphere is reasonably compelling that there are efficiencies to be gained there.

Regina: And I see Rylin has a hand up.

Rylin: I just wanted to jump in about what we're saying in the intern space because it's a really nice example of seeing some things that have been troubling to the public in terms of AI being

used for claims determination in the Medicare space and then that leaving to some government guidance about, sort of what I referenced earlier that the existing rules apply. That if you are unable to deny somebody a benefit based on the merits, you can't deny them that benefit using AI. So that sort of guidance of its use, and how do we think about it in terms of challenges that people may experience and what are the safeguards, I think is an important piece to hold in these conversations.

Regina: That's a really great point. I want to go back to something that Daniel just said, and just kind of probe a little bit around access to legal services. So, you know, you mentioned that your point of view, that there are some things that need to be safeguarded from an ethical standpoint. Lawyers have ethical obligations and zealous representation of clients. But, from, looking out at a nation that is largely lacking access to counsel, what is your point of view about the use of AI in giving more people representation and representatives in filing administrative claims? In this case, Social Security, but there could be many other applications as well.

Daniel: You know, I think there's great potential. I think there's great potential. Right now, I don't think there's been good evidence showing that if you just hand over AI to people who are not legally trained, that they are going to be able to produce documents that are, you know, as high quality as you get from a lawyer. But what the evidence I think suggests is that for certain people who feel comfortable with AI, there's a big potential to help them navigate a space where maybe legal help would be helpful, but not necessary. I think also it may very well change the economic landscape. I mean for right now, it's very hard to afford a lawyer to help you with something like your Social Security claim, right. That's not how, most people I'm guessing that -- I don't have hard numbers on this, that use the Social Security system, they are not using lawyers to help them deal with that, it's just too expensive. But if we have a technology that can allow a

lawyer or someone even like a legal assistant or a paralegal to really perform, you know, good work for you in a fraction of the time, that may change the model. And so, it may end up being that, for instance, Social Security could do something like actually shift a lot more of the claims write-ups and explanations to serve AI support, and then provide some limited form of support to help people say, summarize and emphasize the key points and fill in gaps that exist in their medical records. So, the landscape is shaping a lot, is changing a lot here, and we still have limited empirical evidence. I do think there's reason to believe that the models are changing here and that we may at some point get to, you know, a technology that is frankly, you know, designed to help applicants for Social Security better present their claims and to ask questions. I can easily imagine the start-up focus on that. I can easily imagine the government starting an initiative that is, you know, "Help me file my Social Security claim," that runs you through some questions, that uses a combination of generative AI and structured questions to help prepare first draft of a claim.

Regina: Thank you, Daniel. Question from the audience, how will this affect representative payees? This is a question we've heard more than once from this audience. So, wondering if any of you would be so willing to respond. And I assume that this question is peculiar to the process of filing claims through a representative payee. And something that is very apropos of what you were just speaking about, Daniel, which is a lack of access for many, many Americans to lawyers and representative payees in order to access the program. So, I don't know, if you want to tack on your comments and answer that question too, but if anyone else has any other responses I'm sure the audience would be grateful.

Rylin: This is Rylin. And I was just going to reference that sort of -- it does speak to Daniel's point, that you know, one of the potential powers is to unblock access to information and

navigation of systems for representative payees who often feel incredibly burdened by the process and tracking really critical things in terms of work credits or income, and the worry of having an overpayment. You know, the ability to get the data quickly, to get the data in a way that is understandable to you and empowers you to use it is incredibly helpful. It's also potentially really exciting to think about does it change who feels the need to have a representative payee or at the level of support. Is there ability for SSI to be more customer-friendly and welcoming where you don't feel the need to have that partnership, but you can get that information for yourself and are empowered in a different way? So, I think there's those pieces of -- you know, I've been excited to see people start to experiment and discover what is already possible using ChatGPT or as it exists in Bing where you can ask the question. You know, I have and share some information about types of diagnosis, would that be qualifying for SSI and how do I learn more? That's very different than a few years ago when you were able to ask a chat, you know, help me find the number of my SSI office, and the amount of information that you can get now is different. So, I think there's a piece there that gives us a chance to play a little bit as a representative payee or an individual who is engaging in the system and see what might be beneficial.

Regina: Sorry about that, I was muted. That's really wonderful, Rylin, and I will just go to the last question that we've received from the audience, and this is actually a theme in many audience questions, so I think you can help us by answering this which is one of language access. This person asks, understanding that federal guidance and best practices strongly dissuade entities from relying solely on machine translations for language access, i.e., without any human reviewer, what solutions do you have for AI to increase language accessibility for those who are limited English proficient?

Rylin: This is such a, kind of exciting question, because it is an example of where the technology is shifting dramatically. And we all may, having seen it in this meeting or other meetings we're at. And I actually don't know for sure whether this meeting uses a CART transcriber or uses AI captions. But a couple of years ago, most people would talk about the level of error that happened when you were using an AI system. And now we are starting to see how technology has evolved. And it was about six months ago when I first heard from people, the complaints about the level of error of a human captioner was higher, because we're humans and we, I'll speak for myself, I make a lot of errors. And so that reality piece, and that shift of we get to some data about the, the AI being less error-filled than a human transcriber. Similar evidence exists in language. But it really does speak to that piece about sensitive use cases, and that is how we talk about it, that, you know, if you are using a language translation app, if you are traveling or in a neighborhood where your language of origin is not what is being used at the grocery store, then they'll sort of risk. I might get the wrong coffee order, but the level of harm, not that high. But if I don't have good translation around access to services, and that harmed my ability to receive benefits that I was eligible for, that is a significant level of harm. So it speaks to a need to have those safety checks about translation. That also absolutely should exist in human translation of information when it is about access to benefits and service. Because errors exist in both. And we are now at a place where the data is showing us different sort of rates of errors. And I think that is going to create a shift in that previous, really well-founded bias that it has to be human translation, but now we are in a place where the, the evidence is showing us that there may be different options. But again, when it is a high-risk situation, we need to have that check. And make sure that people are getting the information they need, and that the information is being conveyed accurately.

Regina: Thank you. And I spoke in error, because I have just received two additional questions, that I think would add real value here. So, if you will indulge me, we have a question which I think is very relevant to today's conversation about the disability community. What is the potential for AI to streamline the SSA's SSDI decision making process to help reduce workload and backlogs? So, this is a question about SSDI and backlogs and wondering if any of you would like to comment on that.

Daniel: Well, I will say I think it is very high. I think that at the end of the day, in particular certain skills that we know that GPT type systems are very good at, that haven't been fully exploited. That's the sort of moment we are in, where we know this technology is there, we even have some evidence, but we know it is not being used or exploited by Social Security right now in that space. Some might think document is a key one, right, by getting, summarizing medical records which we have talked about. Writing explanations is a key one. Identifying potentially, potential fraud or questions is a key one. Those are some of the key things that I think that we need, needs to be done in the context of Social Security disability. And at the end of the day, all of them can be substantially improved in efficiency, I think, with the technology that we have. Now, there are obviously risks that come along with it, and we have been talking about a lot of them. And so, the question is, how do we balance that? And in particular, how do we balance that in the public setting versus the private setting? I do think there are differences there, right, so I think it is probably neither surprising nor completely inappropriate for some more aggressive uses of this technology to exist, say in the private sphere, for private incurs, than in the public sphere. But at the same time, I think that there can be a tendency for government agencies to be slow, right. And for changing government to be slow, especially when it comes to new technology is to, and I think that that could conceivably come at a, the cost of actually people

having their claims adjudicated fairly and quickly. And that, that backlog continuing. So, it is a hard balance. I don't purport to know exactly what the right timeline is, but I do think that there needs to be both some important healthy skepticism of all of the risks that are proposed by AI, but also some, some legitimate enthusiasm about the capacity of this technology to make the processing of claims in particular much more efficient and effective.

Regina: Thank you. And so, this next question will have a few parts, but I think it is worth reading out loud. It is very interesting. How can the integration of specialized rules engines, those within government or in NGOs like Urban Institute, and PolicyEngine, with large language models enhance the accuracy and personalization of benefits assessments in the Social Security landscape? That is the first part. Do any folks have responses to that?

Rylin: I like that question because it really speaks to design, and what do we include in setting up the structure that helps guide AI in these use cases. And so it is, it is an interesting question to ask, is do we only include, you know, the legislative guidance in rule making over time in the structure, or do we include some of the excellent work that has happened in the not for profit and other settings that provide insight into the process and the decision making? And, you know, that is a really important question, and it sort of speaks in my mind to the last question about SSDI, because I am shamelessly biased in that I think we should start first with making sure that we're meeting the needs of the disability community. Because if we do that well, and we design great accessible interactive products, then it works for everyone. And move forward in a different way. But my sort of "but" with that, is that in order to do that well, we need to have a wide range of voices, supporting the agency in developing what the structure looks like, and raising questions like the one that was just referenced. You know, are the tools using only the data that comes from regulation? Are they using additional data that comes from court cases over time, or other

factors? And I think that is a really important place for the community and individuals impacted to join and share their perspective at this point.

Regina: Does anyone -- I think Jason might have wanted to weigh in here.

Jason: Yeah, just very briefly. It is interesting to keep track of the difference between expert systems that are explicitly following a particular logic that has been hardcoded into the system. That was fairly common in the previous generation of AI technology. Versus fine tuning a general-purpose model to be biased in a particular direction, and this is the good kind of bias, where it is biased toward following the norms and the cultural rules that you have in your nonprofit or in your government agency. So, they both involve this kind of logic engine, but they are applied in different ways. The older version would sort of give you a deterministic outcome, based on which rules were followed. And the newer versions will sort of steer in the general direction of the rules, but not necessarily guaranteed to follow those rules which has both costs and benefits.

Adler: Yeah, I really like that question as well. Sorry. I realized I was just sitting and listening to everyone else. But I wasn't actually responding. So, my friend Max is one of the cofounders for PolicyEngine, and I really like that organization, because there is a focus on, and I mentioned this earlier when we were talking about the context of helping representatives understand the impact of technology in their day-to-day use, and how that makes their lives easier. And how that can potentially help the people who they are serving. I think with these, these types of nonprofit initiatives, it helps the public better understand how the policies work and how they impact them on an individual level. And they can kind of tailor the interaction to really see what is, what is in it for them. How it impacts them. So, in terms of like policy engagement, I think it is a really powerful tool. But it also helps them advocate for themselves better, because they are a lot

clearer on these really complex subjects how they apply to them. And by the same token, I think it paves the way that people can engage with government when they are asking for support or requesting assistance. It can give them a really personalized view of how these systems impact their day-to-day life and what types of things to ask for. And there was a comment earlier about bias, and it just made me think about the fact that, so the systems can add bias, but it can also detect bias. So I think about a study that a colleague did at work, and she was looking at the way that stigmatizing language was used in electronic health records. So, this is language input by doctors when they are talking about patients, and so instead of kind of summarizing what a patient was saying, they might use the exact words, to kind of demonstrate that it was a non-native English speaker or maybe it was someone from a different socioeconomic background. And then that impacted the type of pain medication they got and lots of other things. And so regardless of the intention of that, the, you know, the study was pretty clear that this is an issue on the human side, and so that could be an interesting thing, an interesting role for AI. Maybe to kind of analyze these types of situations in claims data, and then help advise humans on how to create applications systems that make that, those types of possibilities less prevalent.

Regina: And so, I want to summarize, we -- about a handful of these questions are asking essentially the same thing. And I'll do my best to summarize what there is a theme emerging from the questions that we have just received by email, which is essentially this. Is the AI going to be used in the future as an engine to make determinations? This is sort of a deterministic approach, or will it be dispositive of the agency's determinations? And in doing so, will it use as a data set all past determinations? I think that that question was asked with some judgment that there might be bias introduced by doing so or there might not. So that is one thing to weigh in on. Or other questioners are asking, can it be used in a way that aids reducing administrative

bureaucracies and our main decision makers with more complete information more quickly in order to address more case files sooner? So, I think it would be actually of great benefit to the listener to hear from each of you on this issue. I think there is a lot of questions about what would be fed to the AI in that deterministic sort of dispositive judgment scenario, and what would be fed to the informed choice of, of the Social Security Administration reviewers if they were a human in the loop using AI? So we should, we should go in any particular order. Go ahead.

Daniel: All right, I'll start. So, I don't think there is a world any time soon where a machine is going to make good benefits determination and there would not be any resort to a human appeal of some type. And I don't think there is even any time soon where a machine would make a decision without a human in the loop. But the difficulty really is in a few things. The difficulty is in, is there a risk that the human in the loop is really performing that independent embedded process is really making sure that potential biases in the AI or potential blind spots that it's not detecting are addressed? Or are they sort of just, you know, essentially, you know, sort of green lighting the sort of judgments and recommendations and explanations that are produced by AI? That is a huge concern. And I think that it is something that, you know, I don't have a clear answer to. But I think one needs to pay attention to in designing the appropriate system. I think another element of this that is important to recognize is there is a real risk that when someone gets a lengthy explanation or an explanation that is given by a computer, they don't understand that they have a resort to an appeal. That they believe that it is, well, that is what the machine said, and so we are done. And so I think that even if there is a, you know, we get to a point where these claims determinations are pretty highly automated, the difficulty is going to be in communicating, again being transparent, about look, what part of this was done by a machine,

what part of it was done by a human, what are your rights on an appeal? What is the role of a human on appeal? So, I think those are some of things that I think about in terms of the risks here. In terms of what types of training data, it would be trained on, you know, it is hard to say. There are a lot of different choices that could be made there. But, in general, the way that these systems are operating right now, probably what you would have, is you would have some general-purpose large language model, something like Claude or you know GPT-4, and you would be using that with some retrievable augmented generation, perhaps with something that is known as fine-tuning. So, with fine-tuning, you sort of just use the base large language model and then you adjust it a bit based on your particular data. I don't think it is likely we are going to see foundational models that are sort of trained on just this data any time soon. So, my guess is at least in the short-term, we are not going to be thinking about using Social Security records to train these models. Instead, what we are going to be thinking of is using models with certain additions, Retrieval Augmented Generation being one, maybe some sort of automated prompting, maybe a little bit of fine-tuning to help facilitate the claims process. At least when we are talking about generative AI.

Adler: A transparency piece, that was different, something I was talking about is also with the transparency and what data is being used and where it is coming from. And transparency around when you are interacting with AI versus a human. I think in terms of for the, for the representative, for them to be aware of that and then also for the, for the customer to be aware of when that is happening. And the training data, Dan made a point earlier about the fact that humans make mistakes too, and so I think we would want to approach it from the idea that AI can and will make mistakes. And so, we are not locking ourselves into a situation where AI is the end all, be all. But instead, it's a tool, or it's a block, or it's a part of the process and that is really

clear to people that this is just a part of the process, you can view, or you can get more information on where the data came from. I kind of like the idea that it could use training data, but then give you like specific cases and several of them that kind of help make the point as opposed to, you know, maybe me trying to remember a few that I have seen or thinking about other things that could help contextualize it a lot better. And then the emotional intelligence, we talked about that a lot earlier too. But just in the way that it responds to the, either the representative that is using it or the person and it happened in an emotionally intelligent way so that it's not just kind of like a stop response, but it is personalized for the situation that it is responding to.

Regina: And I think Jason was going to make a comment, and maybe Jason you'd address the part of the question that asked about using past determinations. Past is furlough, past determinations is informing future determinations on a large language model.

Jason: Yeah, if you are not going to use past data, then the question is, what are you going to use to train the model? Are you going to, for example, provide reinforcement learning for human feedback where you hire some people to go over, maybe some present cases or some hypothetical cases and really dig into those cases and say, yes, we think this case should come out this way. This case should come out the other way. And who's designing those case studies for the feedback? And who is being hired to decide you know what the, what the results should be. I understand there is a significant overturn on appeal rate for some of these cases, right. So, if you bring them up on appeal, then the appeal board might come to a different decision. But if you are not confident that you will get the right answer the first time, you know, 99 -- 99 and a half percent of the time, then how can you use that data to train, you know, the AI and what it should be doing in various cases? I also want to touch on the problem of interpretability. Right.

So, you know, Adler mentioned the need for transparency and that is important, but transparency often refers to not keeping secrets. Like let people know that you are using the AI, let people know how the AI is trained. That is very important, so that the affected communities can weigh in on whether the process is fair to them. But even when an agency or a company discloses every scrap of information that they have about an AI's function, they often don't have a true explanation for why the AI is making the decisions that it does. The AI's behavior is often a black box, it is fundamentally uninterpretable. And so, until we solve that problem, until the field of interpretability research advances to the point where we can say, oh, here is why the AI made this recommendation on your claim, I think the more responsible procedure is to have a checklist with the elements that we need to satisfy it. Summarize the evidence for the benefits of the representatives. And then have the representatives weigh in on does this summary show that this particular element in the checklist was met? Because that involves then the decision-making process in a way that activates a sense of responsibility and deters them from being tempted to just sort of rubber stamp the AI's recommendations.

Regina: Thank you for that response. We did just get another question which says, and I believe the IRS is beginning to use AI in attempts to identify tax fraud. Is there any applicability to such tools or initiatives to using AI in fraud detection here? And in such instances would there be data sharing between agencies? So, the question is largely about how soon the government cross agency begins efforts at modernization and incorporation of these tools, and how there will be sharing of information across agencies. I am not sure if any of you are in a particular position to answer that question. But perhaps you can opine about the good use of AI in fraud detection in intervention.

Rylin: This is Rylin. I will start with sort of the sharing piece, because it was referenced earlier in terms of Electronic Medical Records and that the challenge that currently exists that, you know, where information is coded in different proprietary systems, not super delightful to share with each other. And that AI gives us an opportunity for that information to move forward. So in some ways it is a potential incredibly helpful opportunity for individuals who may in fact be eligible for multiple dimensions of Social Services in our country. Where currently, you know, the data and application process is often very siloed. And so would there be an opportunity for things to be shared across in such a way, you know, there could be an opportunity that, you know, a tax filing that shows earned income tax credits and children in a household and other things may trigger helpful information about, you may be eligible for, and would you like us to start sharing information you have shared, so the application process can move forward? You know, there's opportunities to move very siloed information around in new ways with AI. It is a question that, you know, is currently on the table in terms of how government, different parts of government are responding to the planning, review, and process related to the Executive Order. And I think it is an important piece in that. You know, there absolutely we reference several times sort of the fraud determination possibilities of AI. And while that gets a lot of attention, and it is possible, there is that tension that exists in the community about let's be very careful that it is actually finding fraud and not inappropriately identifying someone who may have a very unique circumstance and whose cultural background shows up in a different way, and we certainly don't want to accelerate those biases in mislabeling things as fraud. So those are all considerations of the process. And I feel like we have almost themed words that have come out of this conversation. And for me, they are transparency, that importance of everyone knowing what is being used and when it is being used. And really that, what part is important, you know,

what is happening. And then that challenge of whether we want to say bias or access or equity, but how do we make sure that it is being implemented and designed in a way that really meets the needs of our country, and the individuals that need to be served by these systems. And it goes back to sort of that foundational challenge, of large language models and that they are trained over the world as it is. And the world as it is, painfully is not perfect, and it has some real challenges in equity and access that we have to be really thoughtful about in the implementation of these technologies. And we hear terms like “red teaming” and really making sure that when you are designing and implementing, that you're fighting against those biases that exist and not amplifying them.

Daniel: Yeah, and I will just hit on one or two quick things. So one thing I will say is, I actually do think there is some real potential for not even just generative AI, but this traditional machine learning AI as a fraud detection tool. But I think we need to be really much more explicit about what the tradeoffs are, and like how we think about them. So like to give a hypothetical, right, what if we deployed AI and it helped us detect twice as much fraud? But it also created false positives for 10% greater population, and disproportionately that population of people who got false positives came from historically disadvantaged groups. How do we think about that? Like, I don't know. Like, I don't know what the right answer is. There are a lot of really good things about that example. There are some really bad things. And I think that one of the things that it does, I think become frustrating in some of these conversations, is an unwillingness to get granular and say, here is the level of false positives we are willing to tolerate. Here is the level of disparate impact we're willing to tolerate for these benefits. And at the end of the day, I think we need to have those difficult conversations in order to make determinations about when and how much we rely on these tools. Because again, they come with both a huge upside and a huge

downside. And particularly in a fraud detection context, the huge upside is detecting fraud which makes the system more expensive and reduces benefits available in the long-term for legitimate recipients. On the other hand, the real risk I think is bias and false positive, people who are falsely flagged for fraud which can be a very traumatic thing, very hard to deal with. And certainly, is not something we want to encourage. So having that really explicit conversation and doing sort of the cost benefit analysis is something that I think is important. And it is not currently being thought of enough.

Regina: And Daniel, do you assume that that is being thought of in the analog world? So is there an assumption that by essence being reviewed and evaluated in a paper application context, and you know when you, you mentioned sort of the sniff test of what would be an approachable level of outcomes. I think you are using the word “disparate impact.” I'm not sure you are using it as a legal term.

Daniel: Yeah.

Regina: Under a particular statute. So, I want to dig a little -- as a recovering lawyer, I want to dig a little deeper of what you were saying.

Daniel: Yeah. No, what I would say is I do think that there are some ways in which we become more concerned about bias, and looking at bias when we are dealing with automated systems then we are dealing with human systems, even though actually because in a way, bias is almost easier to detect. And because we can understand sort of the [inaudible]. Whereas it is very, it can be very hard to identify bias in human systems that is not explicit. So, you know, that is a real question. Right now, for instance, I don't know how much the Social Security Administration does in terms of assessing when or how its fraud detection systems create false positives that have a disparate impact. And by that, I mean, just the first prong of disparate

impact, which is they disproportionately affect in a harmful manner, members of a legally protected group, as compared to members of a non-legally protected group. And that would at least trigger scrutiny, it wouldn't necessarily trigger an outcome. I think that there is a limited tendency to think about that in the human setting, because we sort of, you know, at least outside of settings where there is a potential cause of action. And those are not sort of many settings. There is, there is sort of just an assumption we'll do the best we can. So there is some ways in which I do think that concern about bias with AI can actually prevent us from using AI when actually AI might help produce bias. Right. So that is part of the trade-off as well. So AI can actually, if our baseline is human decision making, we are not looking for bias in human decision making, that might become a real problem. And so there is some really, I guess I will pitch a book by a colleague of mine called "The Equality Machine" that basically argues this is a real problem. There are really prominent people sort of push this idea that AI can help us reduce bias because human decision making is so much more biased, and that we become fixated on the risk of bias and AI without thinking about the bias we are currently facing. Again, I don't have clear answers, but those are very difficult challenges that we need to get to be more explicit about. And when and whether it is appropriate to rely more heavily on AI than we currently do.

Regina: So, it is ten till the hour. I have one final question. This time I'm telling the truth, it is the final question. There is a person who's written in, and they are a legal aid attorney. And they are saying that using an attorney for Social Security disability is expensive and most people don't use them. And as a paralegal, at Legal Aid Society Hawaii, my experience is that the only way for our clients to receive a benefit is by appealing to the hearing level which requires legal representation. Also, the attorney fees are set by Social Security and taken out of backpay, so the client typically doesn't have to pay the representative. It is all done behind the curtain at SSA, by

SSA. Legal Aid Society of Hawaii represents clients for no fee. Very few of our clients are approved at application or reconsideration, but at the ALJ hearing level. So this seems to be a comment about the, both the need for representation, how long and lengthy and somewhat voluminous the process can be for claimants, and how fees are transmitted between representative, client, and the Social Security Administration. Do you see any sense of optimism that, that, that folks can be aided and assisted by AI, and further representation, but also this, at this appeals level where many claims are ultimately resolved? I think that was the point of the questioner.

Rylin: This is Rylin. I will start. You know, I think one of the pieces of potential optimism or frankly a call to action should be that we have better decision early on. You know, because there, the fact that we are dealing with such a high rate of things that get turned over at that level is concerning. And what is most concerning about it, is that we know the vast majority of claimants can't access that level. And so, what does that tell us about how the system is currently working and whether the decisions that are made earlier are the quality decisions we all want them to be? And so, is the ways that AI could summarize medical records, do the [inaudible] to a place of fully understanding a case to make a better earlier decision, an opportunity to remove some of that challenge and create greater equity overall. You know, so I think that is one piece to think about. And then there is certainly other pieces about the representation and the need to continue to have the opportunity for appeals.

Daniel: Yeah. One thing I will say is look, in most systems and I don't think Social Security's system is any different, the vast majority of people who are denied benefits, who are vast majority of people who are denied benefits, don't appeal. And the reason why they don't appeal is, they have a tendency to not want to navigate a really complex system. And they need

assistance for it, and the assistance is expensive. And so, I think that that can be a problem. And I do think that AI offers real potential solutions. So right now, there are literally thousands of startups trying to use generative AI to perform different types of functions, and one could certainly imagine a model where, you know, you have a startup AI, that its whole purpose is to help you file Social Security disability benefits. Or again, we can imagine the government investing in that type of thing or trying to facilitate that. So, I have, I do have actually significant hope that legal services will become more broadly available. I think that there are, there are forces aligning against that. I mean, a lot of lawyers are not interested in having robots do their job, and so I think that another space where we need to be cognizant of, it was mentioned that you need to have a lawyer to go through an appeal. I don't know if that is actually a statutory requirement or not, or if it is a requirement at the Social Security Administration imposes or not frankly, or if it is to serve a de facto requirement, but I would like to see a lot of relaxation about the rules regarding when you need a lawyer, and more opportunity for there to be different types of aids, right. A paralegal might be perfectly suited to help someone file a claim and may not need the intervention of a lawyer, especially with some of these technologies. So, I think as these technologies broaden out, sort of the capabilities of folks, maybe have less training, you know, this is a bigger picture point, it may be appropriate to have different types of lawyers just like we have different types of doctors, right. You know, you could be seen by a nurse practitioner, you could get seen by someone who is just, who is not a nurse practitioner, you could get seen by a generalist, you could be seen by a specialist. We actually don't tend to have those categories a lot, and I think there are a number of folks who think that AI will push and should push towards relaxing the, you know, fixed categories, we have either the lawyer or non-lawyer. So that we can make legal services more broadly accessible.

Regina: So --

Adler: Oh, really.

Regina: Go right ahead. Who was speaking?

Adler: Oh, sorry, I was just going to say that I really liked that, that point about task shifting. We think about that in health care. So, I always make the joke that when I graduated from law school, I realized I wanted to be a sassy TV court judge, I didn't want to be a lawyer. So, I'm going to circle back in my 80's and get a show. But I'm sorry, I don't want to speak too far out of my depth. But I think in terms of health care, I can think of a model where there is a nonprofit that has an AI agent, that a community health worker can take out into rural areas and collect information from the patient, and then that information is populated in the system. And then the clinician can log in and access that information. And then they respond to the system and the community health care worker goes back out with that response to the community member. So, I think in models, we are thinking about task shifting to other types of roles that are maybe not attorneys, and being mindful of unauthorized practice of law in a way that you can provide legal information. Technology could be a great tool for that, and AI could certainly play a role in helping insert the attorney precisely where they are needed. And then, you know, allow other people to, to play a part where it is met.

Regina: So, I want to thank all of you, Rylin, Adler, Daniel, Jason. This has been an amazing experience and rolling dialogue, wonderful insights that you shared with everyone. I'm sure that you have the gratitude of the Social Security Administration for weighing in on comments that can help strengthen the disability program. This has been beneficial, beneficial information coming in from the audience. I want to turn it back over to Brian Peltier, the Acting Deputy CIO to close us out.

Brian Peltier: Thank you, and I appreciate all of the work that you guys put together for this panel. I appreciate the panelists as well. You did a great job. I learned a lot. I really appreciate it. Again, thank you to all of our participants on the call. This was a great discussion; I learned a lot. And I kind of want to recap some stuff, but due to the time, we have got a minute left. So, I'm just going to go with the time remaining. But before we close today's forum, I have a brief, few brief announcements. All of those who registered will receive an email with a link to an evaluation for this forum. We appreciate you taking time to complete the evaluation as it will help us improve and offer topics for future forums. If you have any questions following today's meeting, please reach us at: NationalDisabilityForum@ssa.gov. Thank you again for joining us. Remember to join us for Part Two of this conversation on Wednesday, May 15th. The registration will be the week of April 29th. Please look out for an email. Please stay safe and enjoy the rest of your day. Thank you.