Deloitte.

Solving for Tech Ethics™

Episode 1:
Designing and applying ethical technology

Moderator: Beena Ammanath, US Leader, Technology Trust Ethics, Deloitte Consulting LLP
Guest speaker: Paula Goldman, chief ethical and humane use officer, Salesforce

Beena: Hello, my name is Beena Ammanath. I am your host on Solving for Tech Ethics™. Today, we have Paula Goldman, the first-ever chief ethical and humane use officer at Salesforce. Welcome, Paula. It's great to have you on today's show.

Paula: Thank you. I am so excited to be here.

Beena: Paula, you have an amazing background, and I am fascinated by how you got to where you are today. You did your undergrad in political science, then you did your master's in public affairs and then

you've done your PhD, and then you have worked all over the world in different roles. Can you share a little bit about your background and your journey so far?

Paula: Yeah, absolutely. I would say I was very deeply influenced by that first wave of tech optimism. Dating myself a little bit, but like the Internet 1.0 and the boom of just thinking that tech was going to democratize everything. At that time, I was working and living all over the world and helping with startups that use technology to serve underserved populations. So, there

was this idea that not only mainstream technology is going to be transformative, but also this notion that you could specifically use technology as a force to open up opportunities for people. I did that for a long time. I also engaged in impact investing or purpose-driven investing in these startups. At a certain point, I would say, there was like kind of the opposite side of the equation that also started to become clear as technology companies are no longer the underdogs as they were the biggest companies in major stock indices and whatnot. And as the impact of technology

became so ubiquitous in our day-to-day lives, it just became clear that there were a set of very difficult questions that needed to be grappled with and a set of responsibilities that needed to be taken very seriously. And that, just as important as thinking through how you can channel tech for good, you also needed to be thinking about how to avoid, call it what you will, unintended consequences or harms to vulnerable populations and that was going to be a very, very important set of concepts and norms and whatnot, and that's essentially how I got into tech ethics. I was working with Pierre Omidyar, the guy that started eBay and started a tech ethics practice with Omidyar Network. And that was ultimately—we were working with a number of companies—it's ultimately how I ended up leaping over to Salesforce.

Beena: Now you hold one of those unique roles of chief ethics and humane use officer. What does your day look like? What do you cover in a day? It is such a fascinating title and such a fascinating topic.

Paula: Yeah, it's extremely varied, and there is just kind of no shortage of very interesting, complicated, and new questions to think through sometimes. I will say it's a combination. So, my team spans a few responsibilities. I have policy, so policy for how customers use our product from an ethics perspective. I also have ethics by design. So working with product teams on thinking through the life cycle of the products that they are building and making sure that we are building in safeguards in advance and thinking through how we steer the design towards maximum benefit. And then I also recently took on product accessibility and inclusive design, which is both about making sure that our products are accessible to people with disabilities, but also thinking through how when you center a design process around excluded groups, it can lead to so much innovation, and I could go on and on about that. But as you can see,

it's kind of multifaceted. I am literally working with every corner of the business. It's sometimes dealing with customers that have questions about responsible AI, sometimes it's dealing with, like today, the EU released a new draft of regulations around AIs, so we are going deep on that and trying to understand what that means. It's working with product teams as they are looking at product requirement documents for an upcoming release. It's thinking through what policies we should be developing. So it's very, very multifaceted. I think at the base of all of it, it's about listening. It's about listening to our stakeholders internally and equally, if not more important, our stakeholders externally.

Beena: How much of this do you lean into your past experiences that you are bringing? I am really intrigued by the work you have done serving underrepresented minorities and underrepresented communities. How much of that is on top of mind for you as you are looking across all these different teams and products and design? How much of that shapes your thinking and approach?

Paula: When we think about marginalized populations, we think about the impact of our product, that it is accessible to everybody, that it is not harming vulnerable groups. I actually think that's the most important questions we can ask ourselves. Those dynamics vary from country to country and culture to culture. It's hard to think about a more important set of questions to ask. That means that we think about lots of different questions all at once. Last year, in the wake of the racial justice crisis in the United States, we spun up an inclusive language program within product that very much builds on existing inclusive language work across Salesforce but gets very deep into our technical content and code and remediates terms like master and slave and whitelist and blacklist. But we did that knowing that there is kind of the outcome of making sure that we

are not using terms that are hurtful and have a racist legacy, but also that we are thinking about the process and bringing stakeholders into that process. So, we had a number of groups from both internally and externally weigh in and to kind of like what terms need to get prioritized, and who and sort of how we prioritize them and by what criteria and whatnot. In some sense I think in this work, the process is as important as the outcome. The who's involved is as important as the what. And I think that's part of the broader set of challenges with tech and diversity in a nutshell.

Beena: There are a lot of companies who are thinking about ethical tech. I saw Salesforce recently released some of the lessons learned through your journey. Can you touch on some of those lessons that you have learned, some of the best practices on how to think about ethical tech and how do you operationalize it?

Paula: I think what you are referring to is we recently put out a post that talks about what we learned in the effort to develop ethical technology response to the pandemic. To give some context there, we were thrown into a situation where we were quickly developing. First, we developed a suite of tools called work.com, which is meant to help businesses and employers reopen safely. Things like attestations that you haven't had COVID symptoms if you are entering a building, that kind of thing. Then we also created a vaccine cloud, which is a set of tools that helped, at first, public agencies and increasingly a broad variety of customers roll out a vaccine campaign. With all of that, we knew that there were a number of sensitive issues that we had to think through at the intersection of equity and privacy. Because we are talking about sensitive personal health data and because we are talking about a public health crisis that affects everyone, but especially communities of color, in a US context, and, thinking globally, it also affects poor

communities and whatnot. So, having said all of that, and because there is no standard for what does a responsible pandemic response technology look like, we had to start with the basics. First, not reinvent the wheel because while certainly this is the first pandemic that I have lived through, basic privacy principles still apply. Things like collect only the data that you need. Make sure you have a plan to delete it. That kind of thing. Second, you ground yourself in the most vulnerable populations. Again, always essential but doubly essential when you are talking about a big public crisis. When we think about accessibility, what we typically think about is can it be used by someone that has lost their eyesight or lost their hearing, and that is really important when we are talking about people signing up for vaccine shots. It's also really important to think about, like, can you access technology if you don't have a smartphone? Because there is still a segment of society that does not have a smartphone or cannot afford it. It's those types of questions that we really try to think through. So, we approach it with a sense of humility. We don't have all the answers, but we are going to try to share what we think is best practice so that we can learn together with our customers and our stakeholders and then reiterate and revise and learn and keep strengthening the approach. We attempt to be transparent about it and then just to keep learning.

Beena: We hear a lot about privacy in the context of whether it is data sharing for new drug discovery, but also for contact tracing, and those are really hard, complex topics to address, especially when we are in the midst of the crisis. So, Paula, going a little bit back, how much have the internal processes evolved since you took on this role to accommodate for thinking about ethics and technology up front in the processes?

Paula: Well, I can answer that question in a bunch of ways. I started out working on policy. And what I would say is, I think this

is a muscle that we have built. The most important thing, from my perspective, when I was building out the Office of Ethical and Humane Use was that we were creating channels for everyone. Everyone owns this. It's not just some office in the corner that owns it. And that we were creating channels for people to raise questions and concerns, and they would not be penalized for raising questions or concerns. So, we created confidential channels, we created public channels, lots of different ways, office hours, lots of different ways for employees and others to raise questions. I would say there is still more work to do, but I would say there is a muscle that has been built there. I think we are constantly working on the muscle being able to have these really hard conversations and to be able to listen to lots of different points of view on issues that might be contentious. We have an Ethical Use Advisory Board, which contains both external tech ethics experts but also internal frontline employees, as well as executives, and it's a really extraordinary group. We have a lot of really hard conversations about the most difficult issues. I think that's another muscle that we have built is that it's deliberative by design. It is an infrastructure that is meant to take in as many inputs as possible before we make these really hard decisions. I think we have gained tremendous momentum on the product side, as well. For example, with ethical AI, building model cards for all the models that we put out or thinking about the ethics of particular technologies that we are developing. I think we have more work to do there and there is not a blueprint for how to do it, but I am just incredibly heartened with the progress that we have made thus far. I truly believe it's deeply important for our standing in the market, it's deeply important for us to be able to share this work. It is not perfect. There is no doubt. There is so much work to do and so much work left to be done. I guess what I would say, in summary, is, I think we have developed a lot of muscle. It is still early innings for this general body of

work across tech. I think about it like when the security industry got started, like in the 90s, when there were viruses going around and this notion that you could create teams that would help set standards. There was a lot of skepticism at that time, like how could you possibly predict all the different ways that people might try to hack your system? I would say there is a similar set of obstacles around tech ethics today, but I see it going in the same direction, that is towards standardization, towards a set of norms around what best practice looks like a set of processes, and ultimately, a set of kind of shared standards that can be set across the industry and with civil society.

Beena: So, Paula, you touched on this a little bit about how it's important for the Salesforce brand and the products you are building. What are some of the benefits that you have started seeing since the focus on building ethical tech and embedding ethics into everything that you are doing?

Paula: I would say a few things. First of all, customers are asking for it, especially with respect to AI, but you better believe that our public sector customers around vaccine technology are asking and thinking the same questions about privacy and access and all of that kind of stuff. I would say, one, it helps us be more responsive, but as I was saying, we don't handle it as if we are the ones with all the answers and we are just delivering this to our customers. It actually starts a conversation with our customers because they have questions and they are dealing with their own set of best practices and risk mitigation. So, it is an amazing conversation starter, I will say. More broadly, and it's hard to measure counterfactuals, but there is certainly no shortage of tech companies that are in the news controversially for some perceived or real transgression. Certainly it's hard to directly say this work prevented a negative headline, but I am quite certain that is an additional benefit to it. But what I would say, most of all, is it

is incredibly motivating for employees and for our broader stakeholder community. We have a large community of admins, our broad customer community, our nonprofit community or dot org community, people really care about this stuff. People care about working for and with a company that has strong values, and I think even when we don't have all the answers, the fact that we are willing to take a forward-looking stance on it, I think makes a big difference.

Beena: That is so true, and it's becoming one of those, like every company has its ethics manifesto or integrity manifesto, talking about technology and how it is being used and its impact, it's becoming part of that core value prop of the company, the brand. We know that there are different ethics roles that are evolving, whether it's an individual ethicist role or a C-suite role like yours, but at the end of the day, ethics is still a team sport. When you started on this journey, how did you get your leadership and all the employees to rally behind this notion of going down this path?

Paula: I also want to go back to the question you asked before, so let's go back to that. I will answer this question first, which is how do you rally people around it? Essentially by just opening up the door. By opening the door so that people can ask questions when they have got something that just kind of tingles their sense that something might be complicated and worth looking into. Being transparent about the decisions that we have made, both proactively the decisions that we have made and also the decisions we haven't made from a policy perspective. Telling the stories of how certain features got developed and seeing the opportunity to replicate that. So, for example, one of the things that we saw early on was in our ethical AI work. We started developing features that helped customers know when they were developing their own models when they were introducing variables that might create bias in their

models. So, for example, there might be a lot of valid reasons to have zip code as a variable in a model, but it can also be highly correlated with race, which can introduce bias. So having little warning signals in the actual product. We saw that feature and then we saw its potential to replicate it across our product suite. I think just showing those ethics can seem big and abstract and impossibly complicated, when you show these small little tweaks that make a really big difference, they can open up, the light bulb goes off for people and they see how they can also do similarly within their own realm of influence.

Beena: Fascinating. Did you launch it out as an official learning program or was it more of multiple town halls? I am very curious about this because that's one of the biggest challenges as companies start this journey, how do you bring all your employees along?

Paula: We are continuing to experiment. What I would say is, we have done all of it. We do things like consequent scanning workshops where we work with a scrum team and we talk about the product that's about to be built and what its intended consequences are and what some of its unintended consequences might be and how we might steer the ship towards positive, and that's one way that people get engaged. There are certainly formal courses, there are informal courses, there is a lot of top down as well. So, we ask our leaders to talk about it because that signaling really matters, at the same time as we try to inspire champions from the bottom up. For example, we recently had a cohort participate in a Stanford Executive Education course on tech ethics, and they loved the course and they loved talking to each other about their thoughts on these really hard questions. There is a whole panoply of stuff that we try and we continue to experiment. I think one of the biggest learnings I found, though, is that businesses need to own it. So, the more that you can

have within key business units, like key executives that are brought in and they are the ones asking about it rather than some outside party, that makes a big difference. I did want to go back to your question about what the benefits of ethics are because I feel like I left something out. I kind of answered your question in the vein of what's the business case for ethics, but obviously the other arguably more important piece of it is also like what is the impact on society. And actually, these things go hand in hand. So, I talked about employees being super motivated, customers being super motivated. I talked about avoiding bad PR. I would talk about how it becomes something that customers are asking for, but at the end of the day, people are not going to use products unless they trust them. So, when we think about, for example, vaccine technology, if it doesn't serve everyone, it has failed. So, in some sense it's fundamental. There are a lot of cases or circumstances for which the success of its product and its ability to reach a broad population equitably and respect the rights of the people who are using it, it's fundamental to the success of the product and to its impact in the world.

Beena: That is so true, the broader societal impact. I always think of the question just because you are scientists and they could, they did, without thinking about all the possible implications. And we're almost at that point with technology, instead of just blindly building it, take a step back and see what are the ways this could go wrong and proactively address it. Another question that comes to mind is also we talk about tech and, in most cases, we are talking about it interchangeably with AI. Tech in this case and ethics tends to be about AI because that's where the world is at today and most of the usage is around AI. When you talk about tech, do you think about AI specifically or what are the other technologies that you consider in your scope?

Paula: It includes AI, but it is not exclusive to AI. And I think it makes a lot of sense that a lot of the focus has been about AI. I can tell you, you can do a lot of good or a lot of ill with very simple technology as well. So I think we try to apply this to all of our products. We have AI sprinkled across many of our products, but we are equally thinking about what does great service technology look like and what does great marketing technology look like, irrespective of whether we are just talking about AI. I wish there were more talk about just the ethics of basic technology. I think that would be super helpful.

Beena: So, Paula, there are also a lot of people interested in focusing their career on solving for tech ethics. What would your advice be to those who want to follow your path and focus on ethics and build a career around it? What's your guidance?

Paula: Well, it's still early and I think there are tons of opportunities. What I would say is, as I think about the team that I am building, is that there is generally a combination of skills. So it's policy, plus ethics, or it's product management plus ethics, or it's tech research plus ethics. So it helps, I think, to have sort of the

disciplinary skill, so to speak, plus that passion and deep knowledge about tech ethics issues. But I would also say that it's early enough that it's being shaped and there really are no rules. I think it's early enough that people can dive in and shape it. There are tons of companies that are creating their own AI ethics practices. I also hear it styled as data citizenship or data ethics. There are lots of different flavors of it, and I think just getting your foot in the door now while it's still early, lots of things are possible.

Beena: Yes, and there is also so much happening in terms of research and NGO space. Personally, and for your role, how do you stay on top of all the research that's happening in this space?

Paula: Having an amazing team that's staying on top of it! But, also, I would say, just cultivating a network. We have some extraordinary scholars that are on our advisory board, really listening to the news. It's sometimes hard to pick apart because so much of the news is about the latest controversy of the day, as opposed to, if you really follow the research, you get deep insight into the nature of problems that are appearing, but also the nature of potential

solutions. It's a great question because it's actually in some sense more important than following the controversy of the day.

Beena: Paula, thank you so much for joining us on the show. I really enjoyed our conversation. There are so many great lessons and best practices that you have shared across different dimensions. I am curious to hear if you have any parting words for our audience on Solving for Tech Ethics.

Paula: I think the only last thing I would say is this is everyone's job. Whether you work in a tech company or not, whether you work in government or civil society or you are just interested for some other reason, this is everyone's job. And ultimately these are questions that pass through democratic processes. These are some of the most important public questions of our time. I just encourage everyone no matter where they sit to think about the impact that they can have on these issues.

Beena: Thank you for listening to Solving for Tech Ethics from Deloitte. Don't forget to subscribe and check out more on Trustworthy and Ethical Tech at deloitte.com/TrustworthyandEthicalTech.