

**Operator:** Welcome, everyone. Thank you for standing by for the Alphabet third quarter 2023 ESG conference call. At this time, all participants are in a listen-only mode. I would now like to hand the conference over to your speaker today, Ellen West, Head of Investor Relations. Please go ahead.

**Ellen West, VP, Head of Investor Relations:** Hi, everyone. Thanks for joining our ESG call today. I'm Ellen West, the Head of Investor Relations for Alphabet. With me on the call today are Kathryn Hall, our VP of Corporate Legal and Assistant Corporate Secretary and Lead on Governance, and Nathalia Valdes, who leads ESG Engagement for Alphabet's Investor Relations team.

We're also pleased to be joined today by our subject matter expert, Jen Gennai, our Director of Responsible Innovation.

This is the 13th call that we've hosted focused on Environmental, Social, and Governance topics. And we'd like to thank all of you who took the time to send us questions ahead of today's call. Today we're going to address the most frequently asked questions that you shared with us, and in particular, we'd like to take advantage of having Jen and Kathryn on the call with us to address AI Responsibility, and Legal and Governance topics.

I'll now turn the call over to Jen Gennai, our Director of Responsible Innovation. She'll just walk you through a bit of background on herself and how she works across Alphabet and then we'll turn to your questions. Over to you, Jen.

**Jen Gennai, Director, Responsible Innovation:** Thanks, Ellen. Hello, everyone. It is great to have the opportunity to discuss with our investors our approach to AI responsibility.

As the dawn of Generative AI offers us the opportunity to proactively guide the development of an unprecedented technology with responsible practices, it's a very exciting moment to be working on operationalizing Responsible AI and making it real for not just our employees, but all of our global users.

I've been with Google for nearly 17 years, and this year, with the speed and promise of Generative AI, has raised some of the greatest opportunities and greatest challenges of my working career, given the space of Responsible AI is so new and addressing some of the most emerging and complex issues of our time.

I've worked in many different countries, starting my career in my home country of Ireland, but also a number of different teams. But I've always focused on protecting users, ensuring that they're safe online, and managing risk, risk to our users and risk to society.

My first effort in ensuring digital fairness for Google's products was 15 years ago, when I established Google's "Treat Customers Fairly" compliance program, and eight years later, I founded part of the company's ML Fairness efforts. This directly led to establishing the foundation for Google's AI Principles, which were launched in 2018.

These principles that were launched in 2018 are essentially Google's Ethical Charter for the deployment and development of AI. These principles have two main parts. There are seven commitments on how we will use AI responsibly, and what we believe AI should do. Specifically, we try and ensure that our AI is socially beneficial in addressing some of the greatest challenges as well as opportunities of our time.

The second part is a collection of four areas that we've determined we will not pursue. These are what we call our applications we will not pursue because of potential risks and harms. As you'd know, principles are only effective once we put them into practice, so as I was co-authoring the AI Principles with a group of diverse people across Google, I was also designing the governance, practices, processes, and structure that would go along with those principles to ensure that we're putting them into practice, and ensuring that company-wide, we're meeting these stated commitments.

So this led to founding the Responsible Innovation team in 2018, the team I currently lead. This was to operationalize the AI Principles across all of Google. What our team does is we provide oversight to other Responsible AI efforts in the company, as well as defining emerging best practices of what Responsible AI means, and how we meet those obligations day-to-day.

Part of that involves ethical reviews. These are ethical assessments of how our products will operate in the real world. Our ethical assessments identify not just the greatest potential benefits of these AI products, but also the potential risks and harms, and then we work to mitigate those harms or prevent them completely.

At the same time, my team also works on educational content and best practices to ensure that we're sharing this knowledge, not just internally in Google, but also externally in a series of either research papers, blog posts, or otherwise, to ensure that we're up-leveling the industry in understanding Responsible AI best practices also.

Last year, I moved the Responsible Innovation team into our Office of Compliance & Integrity, which indicated the maturity and stage of evolution that we see ourselves at, and especially because of the upcoming AI regulation in this space.

We've also continued to integrate and expand a lot of our processes as part of our Trust & Safety organization, which has been around for over 20 years, always working on keeping our users safe. Together, we've ensured that across all of our products and all of our Generative AI as well as AI releases, that we're operating responsibly.

As AI is still such a nascent technology, there are still many risks to be discovered and defined. Having strong governance in place is just part of the process. We also continue to conduct industry-leading research in these emerging areas to ensure that we're translating research and emerging issues into new governance practices.

So since we launched those AI Principles in 2018, we've built and tested an industry-leading governance process to align our AI projects across the company with those AI Principles. We center our governance around four main pillars: Culture and education; structures and

processes; tools, techniques, and infrastructure; and external engagements and partnerships. I'll share a little bit of detail for each of those.

So for culture and education, it starts with ensuring all of our Google employees and our executives know what are these AI Principles and how do they operate day-to-day and ensure that we're operating responsibly. This happens through various range of trainings, resources, and workshops.

So we've structured employee program trainings that are defined by function. So what this means is that Responsible AI training, that a product manager undertakes may look different than the trainings that we put our sales teams through. That ensures that it's applicable and specific to the roles of our Google employees.

For structures and processes, we adopt a risk-based governance approach, inclusive of our ethics analyses as well as executive accountability to ensure that we're focusing on the most consequential issues that could affect our global users and could manifest in our AI products at any given time.

Our risk assessment framework, which aligns with a lot of the upcoming regulation and drafts that we've seen, helps identify, measure, and analyze ethical risks throughout the AI lifecycle. We map those risks to our mitigations, which can be either technical or policy implemented. We also then develop peer standards of risks of where they are more likely to occur and who are they more likely to affect. We have found that risk is not a blanket experience across all of our products or users, so we have to document and analyze where it is more likely to occur.

This draws upon the best practices of Google's cross-company Office of Compliance & Integrity and our existing enterprise risk management processes.

We also ensure that we're engaging experts in various domains to help us better understand those risks and better understand the mitigations that we can apply to those risks.

We also have senior executive forums that help us debate and discuss the most complex emerging issues that we discover in our AI products. As you can imagine, the emergence of Generative AI has raised many new questions about the potential risks and harms, so we bring together diverse and very senior accountable decision makers to help us work through those problems.

In the third bucket of tools, techniques and infrastructure, we conduct industry-leading research to better understand the potential issues, harms, but also opportunities that AI can allow us. So we ensure that we're applying that research to our tools, to address these issues day-to-day. For example, recently, we've added a new version of the learning interpretability tool, which allows developers to better understand what their model is doing, in order to debug it, but also to also understand how decisions are made through that AI model.

We have also established the Monk scale with Dr. Monk, which allows us to better label and understand how products, especially image-based products, work for different skin tones, to ensure that our products work as well for everyone.

We also frequently update our collection of hundreds of Responsible AI research papers, which are publicly available at [research.google](https://research.google).

And then finally our fourth pillar is external engagement, how we engage with external experts to ensure that Google is always holding itself accountable for understanding some experts in different fields, domains, and community leaders, what are the potential implications for their communities, what are the potential new areas that we may have less expertise on that we should be listening to other people. It keeps us humble, but it also keeps us ahead of the next potential issue.

So we participate in workshops, in user studies, and conferences to understand that we're listening and hearing and adopting the expertise of other people.

One of our ongoing efforts is called Equitable AI Research Roundtable, which allows us to hear from experts in different diverse communities about the potential implications of our AI on their communities. We've pivoted that this year to specifically focus on Generative AI topics.

We also proactively engage with policymakers and regulators around the world to try and understand what is top of mind for them and to share our learnings. We've learned a lot from our Responsible AI efforts across the years, what's worked and what hasn't worked, and so we've been able to provide that as part of our learnings to inform regulation across the world.

We've also used some of those similar learnings to inform global standards, particularly ISO, IEEE, and ongoing work around the G7.

Related, we also engage with ongoing industry forums. More recently we've identified our participation in the Frontier Model Forum and our participation in the White House Commitments.

And then finally, as part of that engagement, we also engage directly with our users. Our users in our products are able to provide millions of pieces of feedback to try and inform and ensure that our products are continually learning and adapting to their feedback, and ensure that they're constantly addressing any potential issues that are raised by our users.

So, as AI has this potential to address global issues facing everyone, from life sciences to climate change, while at the same time expanding the benefits of existing innovations to people around the world, doing so responsibly is key. We need to ensure that we are addressing identifiable concerns, ML fairness, privacy, and safety, in collaborations with stakeholders across the world's AI ecosystems. We cannot do this alone.

We remain committed to sharing lessons learned and to offering transparency into our oversight of our AI systems and how we the adopt our risk management processes, as well as how we hold our execs accountable to our executive decision-making forums and our AI ethics operations.

**Ellen West, VP, Head of Investor Relations:** Thank you so much, Jen. That was an incredibly helpful and clear explanation of how we've been managing these risks and how we're responding to the emergence of Generative AI within that framework.

I'm now just going to ask you to address some of the most common questions that we've received from the investors who are joining us on the call today. The first question was, how does the company oversee AI responsibility? For example, how frequently are you able to interact with senior leadership and how active is Alphabet's board on this topic?

**Jen Gennai, Director, Responsible Innovation:** Yes. So, certainly there's a lot of oversight and engagement by our senior leadership. We have a number of executive level decision-making forums with participants at the Director, VP, and SVP level, which help us address a number of the most emerging issues of our time. We ensure that these executive level forums represent a variety of functional areas. Specifically, we ensure that we have representatives from our Privacy, Trust & Safety, Security, Responsible Innovation areas, as well as then accountable executives from individual product teams, as well as different technology areas, whether it's Generative AI or AI.

So these councils meet very frequently and increasingly frequently, obviously with the speed of our development of Generative AI products and features that we meet weekly on our policy and product launch readiness discussions, and then biweekly at more thematic questions. So, for example, if there's a cross-coding and emerging question of whether different products are adopting our AI Principles in the same way, we obviously strive for consistency, and these executive decision forums are the ones who are accountable to ensuring that consistency across the whole company. So you can trust that your experience on one Google product will be the same as that experience and thresholds around safety and responsibility on another Google product.

**Ellen West, VP, Head of Investor Relations:** Great. Thank you for that.

It's obvious, you know, from your opening remarks and your biography that you've been with Responsible Innovation for a number of years, and our investors have asked if you could share a little bit more of your perspective on how the organization has evolved and what the current structure looks like across the company.

**Jen Gennai, Director, Responsible Innovation:** So, obviously from a compliance and governance perspective, we try and be as stable and consistent as possible, but this technology moves at such a rapid rate, we also have to be very responsive. So yes, I would definitely consider that our governance programs and structures have been iterated upon and improved over the years.

Example of this is when I was founding Responsible Innovation in 2018, this was the first time that we had a full-time team dedicated to our Ethical AI Principles. But, that was based on the work and time of hundreds of Googlers across the company who had already been working part-time on Responsible AI and ML fairness issues and obviously our existing and large scale Trust & Safety organization, which is where I was before I moved into Responsible Innovation.

So we already had hundreds of people participating in these Responsible AI efforts, and Responsible Innovation was merely the next stage of that maturity by adding rigor, robustness, and a formal structure around that.

Since 2018, we've also added another central organization, which is Marian Croak's org is focused on research and technical tools in Responsible AI. So together, my operations and governance team, with her technical and research team, are able to support product teams across the company.

Thirdly, as mentioned, we've focused a lot on education and training. This allows us to ensure that people across the company are up-leveling their ethical critical thinking skills, so that has increased our capacity and scale of how many people are aware of the AI Principles day-to-day and adopting that. More recently, I mentioned one change that has happened is moving the Responsible Innovation team into the Office of Compliance & Integrity. That's in acknowledgement, not just of the maturity of our program, but also the upcoming AI regulation and getting ahead of that.

But we do feel Responsible Innovation has been essentially internal self-regulation for almost five, six years. But as that formal AI regulation comes from across the globe, we are ensuring that we are preparing our processes, procedures, policies for that. So that is a more recent change in response to that evolution and maturity over time.

**Ellen West, VP, Head of Investor Relations:** Thank you for that. Investors also asked if you could talk a little bit more about how we promote Responsible AI across the organization, and more specifically, describe the process that our product teams have to follow to evaluate alignment of any AI tools ahead of their release. It would be really helpful if you could just walk us through how that works.

**Jen Gennai, Director, Responsible Innovation:** Okay. I'll just break that into the two components of the question. I heard that it's how we promote Responsible AI across the entire organization, and then how do we ensure the teams follow alignment with our AI Principles, regulatory tools, or processes.

So for the first part of how do we promote responsible AI across the entire organization, part of it is education and culture. I won't repeat the parts I've said there by ensuring people feel they have the capacity and tools to do that. We say that we try and equip and empower people to make these ethical decisions day-to-day.

But in addition to that, we also essentially talk about it a lot. You have to ensure that your employees know what you stand for. And as our CEO Sundar has said more recently and on our I/O calls, we have consistently said that we will be bold and responsible. We are saying that we are one of the most innovative companies in the world, we're launching the most innovative Generative AI products, and at the same time, all of those will be both innovative and responsible.

So having that culture enforced at the top and ensuring that every employee knows that is critical to how we promote Responsible AI across the whole org.

Then what does it look like at the team level to align their AI products and tools ahead of their public releases? So, for us, in our team, we work with the product team to understand what is the product that is going to be released, and it is a four-step process. We call it our Intake, Analysis, Adjustment, and Decision.

So the Intake is my ethicists and reviewers are assessing what is this product intended to do, what are the best possible use cases of this individual product.

And then we try and understand what are the potential worst use cases? How could this be misused by bad actors? Or how could something go wrong that could potentially show something that could be harmful, could be offensive or otherwise? We then analyze the scale, severity, and likelihood of those issues happening. We then bring that analysis to the product team to understand what mitigations are already in place to prevent these issues.

And if there are not existing mitigations or tools available, we work with our research and technical teams to develop those tools. We also work with our Trust & Safety colleagues to ensure that when this product goes live, are the guardrails built proactively to ensure that we get ahead of any potential new issues.

And then finally, what we would consider the decision of the process is do we feel we have adequately mitigated against those potential risks and harms to allow this product to be released. So that's our at a very high level, simple, four-step process, but obviously, there's a lot of hard work, analysis, and expertise that's brought to the table to conduct these ethical analyses, to conduct the risk assessment, and really wrestle with that tradeoff of the benefits and potential harms to ensure that we're launching our products in the most responsible ways.

**Ellen West, VP, Head of Investor Relations:** Thank you for that, that explanation of how it works in practice with our product teams.

Our next question is about to what degree are different teams under Google DeepMind already aligned to the AI Principles and other Responsible AI efforts?

**Jen Gennai, Director, Responsible Innovation:** Sure. So not many people know that DeepMind were actually one of the co-authors of Google's AI Principles back in 2017 and 2018 as we were writing them, and they were most critical in areas of our four applications we will not pursue. So from the very beginning of these more formalized efforts, DeepMind have been working on the same set of AI Principles that we have been at Google.

Obviously, more recently, we announced that Google's AI research team, known as Brain, has been brought over together with the DeepMind research teams. So that has allowed us to combine our existing efforts around Responsible AI as well as governance with those existing teams.

My team has already worked closely with the existing fairness, ethics, and responsibility teams at DeepMind long before the merger of the overall research teams. So we've been sharing our practices and identifying opportunities for collaboration. With the merge, we have increased the

frequency of our meetings and knowledge sharing, and we continue to identify those opportunities for collaboration. So for me, from the very start, DeepMind has been both informing our policies and principles as well as collaboratively working on our governance and Responsible AI efforts.

**Ellen West, VP, Head of Investor Relations:** Thanks for that. Can you outline the process that your team uses to assess AI-related risks? You referenced it earlier in talking about the process that product teams follow. But how do you actually do the Risk Assessments?

**Jen Gennai, Director, Responsible Innovation:** Yeah. So to do our risk assessments, they're based off of a foundation of a risk management framework. That risk management framework has been iterated on over the past five years annually. We try and ensure that it still fits our purpose and still is addressing the most critical risks as well as our long tail risks.

So that risk assessment framework is not just based off of our existing enterprise risk management approach at a Google level. It's also informed by socio-technical research. So this is research that's specifically focused on users and society and the interaction of technology with the world.

So we establish that AI risk framework to ensure that it's addressing harms and issues related to users and society, not just business risks. So just to clarify the difference of how we consider risks, our potential risk to the people, not risks to Google's reputation or financial outcomes.

So in order to conduct our risk assessments per product, we use this framework as a baseline. It is not a checklist. That is why we have ethicists and people with backgrounds and expertise and civil rights and human rights working on these risk assessments, because there's a lot of deep ethical analysis and critical thinking skills that must be applied to each individual use case. Each use case from a product is different, and we have to understand how it interacts differently with different users across the globe.

So we identify the potential risks that are applicable to this specific use case and these specific users, and then we do that analysis that I mentioned around scale, severity, and likelihood. And many people are likely to be addressed by a specific risk. How severe is that potential risk or harm? We understand that areas that have implications on education or physical health or financial situation are risky in a real world sense, and will have material harm of users. And that may be different and it's definitely considered very different from a risk perspective from a piece of spam going into your Gmail.

So these are the types of analyses that we do. It's not just that risk of harm is likely. It's also how severe will that harm be.

And then once we have that analyzed and assessed, we discuss about potential tradeoff. We have to understand if there's any slight chance of a risk that we are not able to mitigate fully, is the benefit of releasing this technology outweighed by that potential harm, that we can't address every single risk to 100% degree, but is the overall social benefit greater. So that is what we address as an opportunity cost, that if we have to avoid launching a product, is that worse for the world or better? So that's the stages that we do, and overall risk frameworks, which is



consistent, but then is customized per use case, and then using those risks in a tradeoff between the potential benefits.

**Ellen West, VP, Head of Investor Relations:** Thank you for that explanation. Some of our investors were asking more specifically about what we do to mitigate the risk of potential for misuse of AI by customers and other end users, and really specifically those related to human rights abuses.

**Jen Gennai, Director, Responsible Innovation:** So, for us, the potential misuse of AI by customers is also related to unintentional consequences as well, and these are most strongly captured and enforced in our four applications we will not pursue. Those four applications we will not pursue are the fact that we will not develop or deploy technologies that cause or are likely to cause overall harm. We will not develop technology that could be used in weapons or related technology, or in technology that could be used for surveillance violating international norms, as well as any technologies that could potentially contravene our international law on human rights.

So the questions from the investors addressing that is we've built our human rights considerations into our AI Principles, so they are integrated into our AI risk assessment framework that I was just speaking about.

Part of how we mitigate the risk of abuse is building in guardrails to address those four applications we'll not before our product goes live. So even if a customer attempts to misuse our technology, we've built in existing guardrails to prevent the highest risk, highest likelihood, and most severe issues that could arise.

At the same time, we also build in policy guardrails. So whether it's contractual obligations or a flagging mechanism by our customers, where it allows other people to raise the fact that there may be misuse. So, the two main ways are we try and proactively identify areas of misuse, and prevent it or mitigate it, and then separately, if we don't catch it, we have flagging mechanisms for users to raise that to it. So they're the two main ways that we work with integrating our principles proactively via guardrails as well as using flagging mechanisms to address reactively what misuse could potentially be happening.

**Ellen West, VP, Head of Investor Relations:** Thanks. Another more specific question about the social issues associated with the use of AI. Just sort of how can users trust the content they are seeing and what efforts do we make to protect users from disinformation, misinformation, harmful information?

**Jen Gennai, Director, Responsible Innovation:** So we obviously recognize sometimes the power of new AI tools, current social challenges like misinformation and unfair bias, or offensive content, that how users use this product and may be sharing some of this negative content must be addressed in our process as well.

So this is where our Trust & Safety organization, which is a thousand people strong, has been working for years to ensure that users are kept safe online.

So we've automated tools to identify potential areas of social issues and misuse of content to take that down and prevent misinformation. We have shared and updated our policies around what is allowed on our products and what is not. We have actively used AI and ML itself to address potential issues that are amplified by AI and ML technologies, because they allow for more contextual analysis and mitigation of potential issues than we've been able to do before.

So we've been able to use these new technologies to address the new and emerging issues on our platforms.

So in order for users to trust the content, they also need to ensure that it is consistently trusted across our products. So for that, it goes back to our executive forums, I mentioned that, for consistency across Google, we set policies at a Google level that are customized at a product level, but allow for standard levels of transparency, explainability, and trust across all of our products. So users know that Google has set our policies, that we are adopting these mitigations, and that we are responsive to flagging and feedback that we get once a product is online.

So although we try and do that proactively, we also work with specialists to identify where there are opportunities for improvement. So we've been working with the partnership on AI for years to address some of the most critical and emerging issues in AI, and one of the functions we had was around synthetic media, which is allowing people to understand whether they can trust content if it's been developed by Generative AI, or if it is traditional and real media. So we continue to work on efforts like that to ensure that users feel in control, but also understand the content that they're interacting with.

**Ellen West, VP, Head of Investor Relations:** Thanks, Jen. I want to pivot a bit now, because we had a series of questions just on how we're working on engaging with others outside of Google. And I think the first set of questions is around what specific AI frameworks or standards have we found, have you found particularly helpful? Is there an ideal framework out there from your perspective?

**Jen Gennai, Director, Responsible Innovation:** Well, we are constantly reading and researching what other people are developing in terms of Responsible AI, even as we were developing our own AI Principles, which took a year to develop because we were analyzing what had already been produced, what we felt were stronger and would be best informed and applied to Google. Specifically, we looked at things like the U.N. sustainable development goals, the OECD frameworks, and what our global standards organizations like ISO and IEEE have been developing. We used all of that to define our AI Principles and we continue to look at what other organizations are doing from a governance perspective.

So we definitely try and ensure that we are being humble and keeping our standards high and learning from other organizations. Most frequently, we do work with ISO, the International Organization For Standardization, and here in the U.S., the National Institute of Science and Technology, on their AI management framework. We've learned a lot from other industry players as well as the rigor that these organizations have been developing, and we brought some of that learning into our own risk management processes, which I had explained earlier.

But we also share with them. What's worked for us at this scale and complexity at which Google operates, we've also been able to share what works in different situations and at scale, versus things that need a bit more flexibility or nuance.

We also see some very positive advancements in AI frameworks or standards on a regional basis. For example, Singapore released its model AI governance framework and continues to talk with stakeholders in financial services as part of their initiative. And in India, the Ministry of Electronics and Information Technology is considering a proposal for Responsible AI for all. And so there's lots of different regional efforts. Canada, Brazil, South Korea are also working on AI legislation, but obviously, U.S., UK, and EU are more advanced and in their regulatory preparations.

So, we find essentially all of them helpful, and because they all have different parts that we can learn from, adopt, and even pilot in some states, that if they are not already available, we're able at the scale and as frequency we conduct our reviews to pilot other regions' or other organizations' learnings to see if it actually works for us.

So I wouldn't yet say there's an ideal framework, because this technology is so new. I think it's most important that we're iterating and learning, instead of focusing on have we already achieved the gold standard. The technology is just moving too quickly to say that there is one yet, but we feel we're learning a lot from these other organizations to constantly be holding ourselves accountable.

And obviously, as part of our White House Commitments that we recently announced, as well as our participation in the G7's Code of Conduct for organizations developing advanced AI systems, we are both learning and contributing to those AI frameworks and standards.

**Ellen West, VP, Head of Investor Relations:** And then following on from that, how has our collaboration with others in the industry, and with partners, been evolving over the past year as all of this has been moving at such a fast pace?

**Jen Gennai, Director, Responsible Innovation:** It's allowed us to focus a lot more over the past year. Previously, there were many different partnership efforts that we had across a number of different themes and technologies, whether it's explainability, interpretability and fairness, working with partners across partnership and AI, G7, OECD, all these ones that I've already mentioned. But over the past year, with the increase in attention and scrutiny on Generative AI, there's a lot more focus on the technology of Generative AI and the key unique issues that Generative AI has raised.

So although there's a lot of overlap of issues across AI and Generative AI in areas like fairness and privacy, we've also seen a lot of unique issues around Generative AI, around areas like factuality, misinformation at scale, that are similar topics, but in a very different scale and speed than we've seen before.

So we've created across industry forums the frontier model forum to develop standards and benchmarks around this very specific technology of Generative AI, but one that is moving so quickly, that's allowed us to focus our efforts in these areas.

Similarly, although we've been engaging with partnership on AI for years, the synthetic media framework that they've developed has allowed many industry groups, including ourselves, to develop and foster best practices around synthetic media development, creation, transparency, and sharing.

So we continue to engage there across fairness, privacy, and all these other topics, but really this past year has allowed us to focus on the most critical technology, which is Generative AI and all the unique issues it's been raising.

**Ellen West, VP, Head of Investor Relations:** Thanks. We have all these efforts you described, we're engaging in all of these collaborations and conversations, but how can our ethical AI Principles benefit the smaller players who are using our LLM for their own AI capabilities? You know, our customers and partners who may not have a Responsible Innovation team like yours, how are we sharing what we do with them?

**Jen Gennai, Director, Responsible Innovation:** I would agree, no matter what size a company is, as they're developing LLMs and AI, that their AI must be developed responsibly. But you're exactly right, that not every company has the resources and full-time team like mine, as Alphabet does, to work on these issues. So we also consider that as part of our responsibilities. It's a responsibility to share our knowledge with other players.

So there's a number of different ways that we do that. One of the key ways is sharing our research. We've been publishing our research for years, and not just AI technology and underlying considerations, but also in Responsible AI. I mentioned earlier that we've published hundreds of papers in the Responsible AI space, whether it's in fairness, privacy, security, transparency. We make that available to the research community, so they can use that in ways that are right for different industries, different considerations, different use cases, different regions, and it really democratizes access to the information that we have developed as part of our advanced research.

The second way that we make it available is by sharing our process and governance. My team makes available the practices around our risk assessment framework, our harms technology. We've published those so other people can off-the-shelf adopt our risk framework as well as our harms taxonomy day-to-day, so they don't have to have the skills of what went into developing these very advanced frameworks. They're able to adopt out as well as the governance processes that we share on Google's website, as well as in our annual report.

And then thirdly is the tools and infrastructure. We provide Responsible AI tools via TensorFlow, available for anyone to adopt and deploy in their usage of LLMs, as well as obviously our Cloud commercial department makes available the most advanced Generative AI products that have Responsible AI considerations built in. We have safety guardrails, we have explainability, we have fairness and testing and considerations built into our Cloud commercial offerings. So hitting across research, process and governance, and the technology and tools, we do consider it our responsibility to make those available and externally for free and as part of our commercial offerings.

**Ellen West, VP, Head of Investor Relations:** Thanks for that. You've talked about how your team is structured, our governance and processes, and the recent move into the Office of Compliance & Integrity as we are getting ready for the regulatory environment. Could share a little bit more on our view, Alphabet's view, on the AI regulatory environment and compliance with the pending regulations that are coming.

**Jen Gennai, Director, Responsible Innovation:** Yeah. So we've been obviously -- everything I've talked about today, we've been hard at work for years to develop and deploy AI responsibly. But, it's also -- that's formed our internal regulation, what we would consider internal regulation. We are a team that sits outside of all of our product teams, so we provide that oversight and essentially those policies or rules of what is acceptable or not on our products.

So we have been able as part of this very fast-moving discussion space with AI regulation -- obviously, we haven't seen all of it launched yet, but a lot of discussions, especially after the Generative AI pushes and launches that we've seen, that's how it's increased in both considerations, regional discussions, global conversations around what AI regulations should look like and how to ensure that it is both flexible enough to allow for new technologies like Generative AI, being specific enough to keep users safe and ensure that they can trust our technologies.

So, for us and as our CEO has said, AI is too important not to regulate, and it's too important not to regulate well. So, for us, that means that regulation is critical. Especially critical in emerging spaces where we do not yet understand all the potential risks.

But it also has to allow for all the potential social benefits. I mentioned earlier that AI and Generative AI potentially can help us address issues at a global scale that have never been able to address before, that it can allow us to address emerging technology issues, it can help us address climate change, mitigations, life sciences. So we have to ensure that we're supporting those socially benefit outcomes while mitigating the worst of the risks.

From that AI regulatory advancements that we're seeing, obviously part of our expectations and what we believe will lead to greater trust is about more interoperability. Instead of seeing so much difference in fragmentation on a regional basis, we know that there are common themes that users expect and will keep them safe. Themes around explainability, fairness, privacy. These are known things that we all must ensure that the industry is up-leveling and ensuring safety in, while allowing for that technology to be advanced in the most critical way.

So a kind of call to action as we see these regulations and what we've shared as part of those discussions is this notion of interoperability as well as risk-based approach. Not every use case, not every product is equally risky for our users. So we need to take that into account. So we have been participating in sharing that based on our learnings in our own products.

**Ellen West, VP, Head of Investor Relations:** Thank you. And our last question is a bit of a shift in focus. It's about AI talent. So, two different questions. Do you see AI talent shortages, and/or diversity challenges to the adoption of AI, and responsible adoption of AI?

**Jen Gennai, Director, Responsible Innovation:** I'll start with the latter. Specifically with the adoption of AI, I certainly see diversity challenges for that. How we consider it is, if an AI product or technology is not built for everyone, it obviously negatively affects that technology's engagement, trust, reputation, and ability to grow. So in Google, we've obviously kept it critical, not just to build products that are inclusive of everyone, but to ensure that we're hiring diverse teams, that they are the ones who can really say whether our product will, in fact, work for everyone.

Part of our risk assessment framework is to consider who will benefit from this technology. Who could potentially be excluded from this technology. And also, who could engage with the technology, but may not have a similarly positive experience as other people. So we definitely see that diversity of talents, informing AI allows for the diversity of experience, that's what happens afterwards.

For us, our ML fairness efforts, which we've been conducting since at least 2014, specifically to ensure that AI does not create or reinforce unfair biases, and that it doesn't expose or amplify unfair stereotypes, that we need to ensure that diversity of talent is reflected in the diversity of tools and accessibility of the products that we deploy.

In terms of AI talent shortages more broadly, in the Responsible AI space, we are thankfully seeing a lot more universities and technical courses considering ethics and Responsible AI as part of their learnings, instead of adding it on as an additional module. There is positive because it is seen as not just adding on ethics as an optional part of the development of technology, but it's a critical and inherent part of the development process, which, again, ties with how our own CEO has described that our advanced Generative AI and other AI efforts will be bold and responsible. They must be tied together.

And then I'll just end on the overall question around AI talent shortages. I think we are aware that our STEM courses and participation needs a lot more people participating in order to support this advanced technology future.

**Ellen West, VP, Head of Investor Relations:** Thank you so much, Jen, for taking the time to so thoughtfully respond to all of these incredibly complex and interrelated questions.

And with that, I think we'll move on quickly to just recap with Kathryn on the kind of top four areas of questions that we received from investors on legal and governance topics. Kathryn, I'll hand it off to you.

**Kathryn Hall, VP, Legal, Head of Corporate:** Thanks so much, Ellen.

Hello, everyone. It's great to have another opportunity to speak with you all about legal and governance matters, and I'd like to thank everyone who submitted questions.

So I'll dive right in. Starting with voting rights. So, investors have asked whether the Board has considered transitioning to a single share class with one-vote per share. We've also been asked to share the Board's continued rationale for maintaining the multi-class structure to date.

So this is a topic that the Board reviews on an ongoing basis, and as a result of that review, they have concluded that they believe that the current capital structure supports our long-term oriented focus while allowing for appropriate accountability. Our success and value are owed in large part to the leadership and vision that was originated by our co-founders, Larry and Sergey, both of whom are still very actively involved in Alphabet.

Our company, as you know, is now led by our CEO, Sundar Pichai, who has been with the company for close to 20 years. This value creation is further enabled by the stability provided by our capital structure, which insulates us from short-term pressures and gives us greater ability to focus on long-term interests than might otherwise be the case. This is evidenced by our company's history of delivering exceptional returns to our stockholders, and the Board believes that our governance structure and independent Board leadership helps ensure effective oversight alongside our capital structure.

Moving on to questions about Board composition. A number of you asked us to address questions regarding our Board composition, and specifically thoughts on director tenure and Board diversity. So, our Board is composed of seasoned leaders and visionary technologists who bring broad knowledge, deep experience, and specialized expertise that enables us to help guide and oversee the company's progress in these areas.

The nominating and governance committee regularly assesses the appropriate size and composition of our Board, the needs of our Board, and the respective committees, and the qualifications of candidates in light of these needs.

We believe that director diversity strengthens the function of the Board, and they do consider diversity across a multitude of factors, including professional experience, education, race and ethnicity, gender, age, and cultural background as part of the Board refreshment approach.

Three of our most recently joined independent directors -- Marty Chavez, Robin Washington, Frances Arnold -- all of whom have joined in the past five years, support the balance between having members who have been on the Board for a very long time and those new members who bring fresh perspectives. So we have that balance of, you know, deep institutional knowledge as well as fresh perspectives.

About 2/3 of our Board comprises directors who self-identify as female or from an underrepresented community, and our Board diversity has trended positively with director refreshment process.

And we're very proud of our Board, and believe that we've developed a Board that has a very high level of expertise and skills with the experience necessary to oversee our highly complex diverse global business.

Moving on to shareholder proposals. We've been asked to provide some insight into how the Board evaluates and addresses shareholder proposals that have received notable shareholder support. So, let me start by just sharing some context around the process that we undertake to evaluate, manage, and consider our responsiveness to the topics raised in the stockholder proposals that we received.

Our management team and our Board takes these -- takes the proposals very seriously, and the topics that are raised are in areas where we are already doing meaningful work, and that we care deeply about, because they support our ability to create long-term value for our stockholders.

Throughout the proxy season, the full Board and relevant committees receive updates on the proposals that we've received as well as our recommended courses of action. And then following the annual meeting, the Board reviews the outcomes of the meeting and engages with management on how to be responsive to the topics of most interest.

And we take a very targeted and thoughtful approach on how we manage these. This approach includes working closely with subject matter experts to not only review and evaluate the asks, but also ensure that our response to the proposal takes into consideration any actions or significant work that we already have under way, which is frequently the case.

We may also engage directly with the proponents to understand better their perspective and their recommendations. We really value the feedback that we get from our investor engagements as well to help inform how we evaluate the programs that we have under way, the actions that we've taken, and the topics that have been raised by some of these proposals.

So, switching gears. The final question that we're going to address today relates to how we're thinking about public transparency as we move forward in the U.S. antitrust trial, and whether we have any comments on the trial. I'll start just by saying this. What the case comes down to fundamentally is this. People don't use Google because they have to. They use it because it's helpful. There are more ways than ever to search for information today. Things like TikTok, ChatGPT, Instagram, Amazon, Reddit, Pinterest, and many, many more.

We work to build innovative tools that help people find what they're looking for. We launch thousands of changes every year to improve Search, and apply technologies like AI to make it even more helpful. Some examples include innovations in the areas of Lens, Mum, and Search Generative Experience.

This case is a backwards-looking case at a time of unprecedented innovation. Look at the recent breakthroughs in AI, the new apps and services, all of which are creating more competition and more options for people than ever before.

So we're going to continue to focus on building great products that help people in their daily lives, and we look forward to continuing to argue our case in court.

**Ellen West, VP, Head of Investor Relations:** Thank you, Kathryn for walking us through those answers to the top questions we've received on legal and governance topics.

I also want to thank all of you who submitted questions that we weren't able to get to today. As I mentioned, our focus on this call was Responsible AI and Governance, but we welcome all of your questions. They do help us plan topics for future calls, and which subject matter experts that Alphabet are best positioned to address to your questions.



One thing I wanted to draw your attention to, because we've received a number of inquiries from all of you, is that we have published our updated environmental report, where we describe efforts and outline how we're driving positive environmental outcomes throughout our business. The report does address how we're developing new ways to make AI computing more efficient, while leveraging the opportunities that AI presents to have a positive environmental impact.

Yesterday, we had a Sustainable with Google event in Europe, and if you just Google it, you can find the video from that, where we talked about more of our innovations on the climate front. The report also discusses our climate-conscious data center cooling strategy, and has a section dedicated to water stewardship.

And finally, an area I know we've talked to many of you about, an overview of our public policy engagement and advocacy efforts around climate change and the environment. Again, just encourage you to review this report. It's a substantial report compared to the brief reports we've issued in prior years. We also make many other ESG materials available on the ESG section of the IR website.

And finally, just want to thank all of you for the ongoing continued constructive feedback and dialogue throughout the year. If you have additional questions or feedback for us, please send an email to [investor-relations@abc.xyz](mailto:investor-relations@abc.xyz).

That concludes our call for today. Thanks, everyone, for joining us.

**Operator:** Thank you, everyone. This concludes today's conference call. Thank you for participating. You may now disconnect.