

# NOW Podcast

SEASON 3 | EPISODE 3 TRANSCRIPT

## Getting Ethical AI Right- And Why It Matters to Us as Investors

Guests: Kay Firth-Butterfield, *Head AI and Machine Learning, World Economic Forum*  
Aza Raskin, *Co-Founder, The Center for Humane Technology*  
Arathi Sethumadhavan, *Head of User Research for Ethics & Society, Microsoft*

Host: Katherine Kroll, Investment Specialist, *Brown Advisory*

### Investment Commentators:

Victoria Avara, *ESG Equity Research Analyst, Brown Advisory*  
John Canning, *CFA, Equity Research Analyst, Brown Advisory*

### Ken Stuzin (00:03):

Hello, this is Ken Stuzin. I'm a partner at Brown Advisory. Welcome to the NOW podcast. NOW stands for Navigating Our World. Through these discussions, we try to understand the world better to navigate some of the most pressing questions that are shaping our lives, our culture, and our investment challenges. As we look to the future, whether we agree or disagree with each other, the one thing we know for sure is that none of us can figure this out on our own. At Brown Advisory, we are focused on raising the future and we hope these NOW conversations will help us do just that.

### Katherine Kroll (00:48):

Right now, as you listen to this, there are computers and network servers all over the world pondering one topic, you. They are improving your life invisibly, streamlining traffic on your commute, helping drive down your utility costs, curating playlists on your favorite streaming app, and likely improving the health care diagnoses you receive. But to do this, they watch your face. The way you walk, how fast you type, they monitor your financial habits, the color of your skin, your social media activity, and then suggest content that might amplify or contort your existing views.

### Katherine Kroll (01:27):

This reality is the basis for one of the most powerful technological revolutions in history, artificial intelligence or AI. AI solutions could contribute \$15 trillion to the global economy in 2030, more than the current economic output of China and India combined. But AI is often positioned as a binary choice, either glorified as our only hope to mitigate impending disasters like climate change or vilified for corroding our social fabric. Of course, the truth is much more complex. I'm Katherine Kroll, and I'm an investment specialist at Brown Advisory where we are enthusiastic investors in AI and the opportunities it informs. We believe it is our job to understand and allocate capital to these opportunities and to understand the potential downside of the risks. Applying an ethical lens to AI positions us to maximize the upside while recognizing that the bottom of AI is still unknown and that we must be diligent in our efforts to learn as much of what is known as possible. That's why we invited three experts on the topic to discuss why getting ethical AI right matters and what might be at stake if we get it wrong. Aza Raskin is the co-founder of the Center for Humane Technology (you might recognize him from the documentary, *The Social Dilemma*). Arathi Sethumadhavan is the head of user research for ethics and society at Microsoft. And Kay Firth-Butterfield is head of AI and machine learning at the World Economic Forum.

**Katherine Kroll (03:07):**

After talking with Aza, Arathi, and Kay, I sat down with two of Brown Advisory's technology analysts to dig into the investment implications. I wanted to start by asking Aza to begin with the basics.

**Aza Raskin (03:22):**

So what is AI? It's sort of a difficult question because in the end, AI is just matrix multiplication. But the way I like to think about it is systematic large-scale trial and error, blind to human values. That is, it's a big black box that is attempting to optimize whatever we tell it to optimize. And it does so without any consideration of morality, ethics, what makes something right or wrong. It just does what's effective. And one of the largest philosophical quandaries over our times, the sort of challenge that humanity faces, is can we learn to do what's right in the face of what's effective?

**Aza Raskin (04:07):**

In AI, there's this parable of the paperclip maximizer. It is if you tell an AI, hey, I want you to make as many paperclips as possible, that's your job, it doesn't realize that tearing apart the world for materials to make paper clips isn't part of its job. And yet it'll happily do that. Really, this is going back to a very old human story or myth of be careful what you wish for, because it'll come true, just not in the way that you intended. So let me give a really specific example of this kind of thing. This actually comes from Guillaume Chaslot's work, who's a YouTube dis- and misinformation researcher.

**Aza Raskin (04:49):**

YouTube, you think fine. It's just maximizing the amount of hours watched. Well, one is one adversarial way that YouTube's algorithm might discover for getting people to spend more time on YouTube. One is finding legitimately more engaging content. Another would be to find content, which says don't trust the mainstream media, because the more you don't trust television or other media sources, the more you turn to YouTube. And in fact, when he was doing his research, he discovered that YouTube was indeed asymmetrically recommending content that's conspiratorial in nature that says don't trust the mainstream media. So this is exactly the case of the paperclip maximizer. We wished for one thing, engagement, we got another thing, destabilizing democracy and our epistemic comments.

**Aza Raskin (05:37):**

Why is this so important? Well, what's sort of funny is that Siri has decided that I was talking to it in between. And I saw an entire running log of everything I just said. AI is taking over.

**Katherine Kroll (05:53):**

Oh gosh, that is poetic.

**Aza Raskin (05:57):**

Why is this so important? Well, 70% of all YouTube's watches comes from AI generated recommendations. What content then is getting the most watch time? Well, the words and titles that are most recommended include obliterates, demolishes, damages, smashes, forces to admit. It is the worst side of humanity that of course grabs our engagement. And so we, as a species, are locked into the worst of who we are. We see the worst of who we are. And so we become the worst of who we are.

**Arathi Sethumadhavan (06:40):**

These machine learning systems, what happens is they are trained on these mountains of historical data that says, hey, these inputs have generated these results in the past. So now let the machine try and figure out what the connection is. And after a period of training, if you are successful, it comes out with a set of rules. So AI systems, they need a lot of data for training and for making these predictions.

**Katherine Kroll (07:05):**

This is Arathi Sethumadhavan, head of user research for ethics and society at Microsoft.

**Arathi Sethumadhavan (07:11):**

So what I'm trying to say here is that it is never going to be perfect. Lot of people think that AI is a like magic, right? And you see it in the movies or something, but AI is actually everywhere. If you come to think about it, you look at your ride sharing apps, like your Uber or your Lyft, they are using AI to determine the price of your ride or to predict ETAs and match you with other passengers to minimize your detour. The spam filters in your email, they're continuously learning from a variety of signals, like the words in the message, the metadata, who has sent it.

**Katherine Kroll (07:49):**

Okay. So AI is everywhere and we're benefiting from it in our everyday lives. When you think about the AI applications at scale in your work at Microsoft, what do you are some of the opportunities for AI to solve our thorniest economic, environmental, and social problems?

**Arathi Sethumadhavan (08:05):**

You can think about it transforming literally any industry that you can think about. Health care, for instance, you can use it for detecting serious health conditions, right? Like malignant skin cancer. It's been utilized a lot in agriculture where things like what is the most optimal time to predict sowing seeds to transportation like your automatic parking and collision avoidance to the retail industry to predicting trends. So it's across the board. But unfortunately, a lot of these technologies are also being developed with very little assessment of the impact that these can have on people.

**Arathi Sethumadhavan (08:41):**

There's an article that I saw recently on how AI enabled cameras are being used to read distress on the faces of women in Northern India. And I have huge problems with that because there's a lot of uncertainty about the validity of a technique like this. What does that even mean to read distress? Right? Last year in the midst of all the Black Lives Matters movement across the nation, we saw an article around how an African American gentleman in Detroit was wrongfully arrested due to a faulty face recognition. There are lots and lots of examples that we see today, right? So clearly AI has got huge impact on societies. So it's very important to design these sort of technologies with intention.

**Aza Raskin (09:29):**

There is a really famous study the *Wall Street Journal* covered where Facebook asked users to rate how good they thought a particular post was for the world. And it turns out users are pretty good at determining good for world, bad for world. And Facebook had the opportunity to down rank content that users think are bad for the world, and yet Facebook chose not to do it because it hurt their bottom line metrics. And I think this is the other really important thing to note about AI is we think when we hear the paperclip maximizing story, that the obvious answer is, as soon as you see AI starting to misbehave, these recommendations systems starting to misbehave, you just turn it off. The problem is, is that it's hooked itself into the core capital logic of our world. And so, now our business models and our bottom lines depend on the AI. So we can't just unhook it. I think the phrase that comes to my mind is that it's not that that AI or technology is an existential threat to humanity. It's that the worst of society is an existential threat to humanity. And technology is amplifying the worst of society. We live in a world where tree is worth more dead than alive, and a whale is worth more dead than alive. And so long as that's true, that the commodified form of the world is more valuable than so too will human beings be worth more isolated and lonely, misinformed, disinformed, outraged, and divided than we are as full living human beings. And I think that actually points at a really interesting solutions, because we live

in a world where it's very hard to hear other people, to see them for who they are, right? We are constantly seeing an infinite stream of evidence for how awful the other side is. Right? Think for Black Lives Matter, the left will see video after video after video of police brutality, it's not wrong. And those on the right will see video after video after video of protestors beating up on police or setting fires to buildings, and each group says, how could the other side possibly believe what they believe? Aren't they seeing what I'm seeing? And the answer is they're not, they're not seeing what you're seeing. So how can we get over that? There's a really interesting group called More In Common, and they did this research called the perception gap research. And what it showed is that really there are two crises happening at the same time. There is the truth crisis. We're all sort of familiar with truth decay.

**Aza Raskin (12:24):**

And then there's another crisis sitting on top of that, that we're a little blind to, and that's the perception crisis, that we are perceiving the other side incorrectly. So an example: If you ask Republicans, what percentage of Democrats are LGBTQ? Republicans will estimate that it's over 30%. In reality, it's less than 6%. And then if you ask Democrats, hey, what percentage of Republicans think all Muslims are bad for the US? Democrats will estimate that's on the order of like 85%, where in reality, it's less than 15%. And this incorrect viewing of the other side just goes on and on and on about our most basic and our most important concepts. So when we think we are fighting with the other side, we're not actually fighting with the other side, we're fighting with a mirage, an image of the other side that is not true. And I think when we look into our phones, we think we're looking into a mirror, but we're not looking into a mirror. We're looking into a fun house mirror, a distortion of who we are, and we become what we think we see in the mirror. So if we were to borrow some of the concepts of nonviolent communication to scale up computational nonviolent communication to a society size, well, what would that look like? Well, the perception gap is actually really easy to measure. Unlike truth, which is really hard to make measure. With the perception gap, I just ask you, what do you think the other side believes? And then you ask the other side, what do they actually believe? And you can measure that difference.

PART 1 OF 4 ENDS [00:14:04]

**Aza Raskin (14:03):**

And then you say, hey, any content which increases that perception gap, let's down-rank it. Let's not show it as much, and any content which helps us accurately see the other side, let's give that further distribution. The realization then is content that polarizes, that creates anger and hate, almost always capitalizes, preys upon a false image of the other side. I think we could start to tackle misinformation, disinformation by solving the perception gap.

**Aza Raskin (14:40):**

Just a couple interesting points more on the perception gap is the more social media you consume and normal media, the greater your perception gap. The more people post on social media, the people that post the most are the ones who have the highest perception gap, so sort of a double whammy. Then, the bigger the perception gap, the greater likelihood you are to view the other side as bigoted, hateful. And so, to me, this kind of thinking points in the direction of how do we create a listening society where it re-knits together all of the disparate sides so that we can move forward.

**Katherine Kroll (15:17):**

I really like how you put that. It reminds me of a quote of yours that listening used to be cheap and speaking very expensive, and now it's flipped and anyone can speak, but not anyone can listen. I'm just curious, how do we get AI to listen in a way that is solution-oriented to the challenges you just outlined, certainly, but in a way that also brings some of these other stakeholders who are benefiting as it stands right now from not listening?

**Aza Raskin (15:50):**

I think we have to look really long and hard at what are the objective functions of our AI, because inevitably what happens is we will make a wish and it'll come true just in a way that we do not expect. So, if we want to have AI that listens, we have to change the relationship that we have to the systems that have a kind of asymmetric power over us.

**Aza Raskin (16:23):**

Let me be a little more specific. What Facebook and Google and Twitter and TikTok can predict about you is strictly more than what your doctor, your therapists, priests, or lawyers can predict about you. Because, given the data of where your phone is, which is of course shared behind the scenes between all the different providers, how quickly you move your hand, how your eye moves around the screen, you can predict when someone is having affair, if they're having an affair, whether they're starting to get dementia, if they're depressed just based on how somebody's hand moves. We did some of that research at Jawbone. If you have a webcam and eye dilation, you can predict whether they're drunk, whether they're on MDMA or another drug. You can predict when somebody has lost their job. You can predict somebody's big five personality traits.

**Aza Raskin (17:13):**

The asymmetric knowledge that these AIs have about us means that they can put us into vulnerable exploitive situations. An example of this happening in the real world was Facebook discovering that, shock horror, teenage girls that are depressed are much more susceptible to ads about makeup or body image. The AI not knowing that that was wrong just finds those vulnerabilities and uses them. I think, in order to move to an AI that listens, it has to be in a fiduciary relationship to us, has to have a duty of care towards us, and the greater of asymmetry of knowledge it has, the greater power it has over us, the greater the degree of liability we need to have for a misalignment of it acting in our best interest.

**Katherine Kroll (18:11):**

Kay Firth-Butterfield is one of the foremost experts in the world on the governance of AI. She believes it's key to tackle the issue of trust.

**Kay Firth-Butterfield (18:21):**

The public doesn't trust AI very much at the moment. It used to be that the public didn't understand it, and then it was that most of the newspaper articles came with Terminator attached to any article on artificial intelligence. But now the public are seeing some of these problems that you find in AI writ large. They know, for example, that if they apply for a loan and it's an algorithm that assesses it, then they might well find themselves turned down because of bias in the data that went into the algorithm. Trust in AI has become a huge topic, not just with the general public, but also with investment, and also with your employees.

**Kay Firth-Butterfield (19:16):**

We've seen some very high-profile problems that some of the tech companies have encountered and the way that their employees have addressed that. On the investment side, we're seeing increasingly now that investors and VCs want to look at investing in companies that have proper ethics codes, have proper ways of being responsible in their use of AI, because without that, they are likely to run in to all sorts of problems around what the product looks like, how it's deployed, what its liability is, and also whether the customers are going to like the deployment of artificial intelligence or whether they'll get a backlash from it. That's even if we take out the fact that there's likely to be regulation around it.

**Katherine Kroll (20:18):**

I'd love to hear about education among consumers and who else needs to learn about what AI is and how we are consenting or not to its use in our everyday life.

**Kay Firth-Butterfield (20:35):**

You need to make sure before you start deploying AI that you have educated your board. We did some work with boards on AI and found that dramatically large numbers of directors of boards who didn't actually understand AI, so therefore couldn't possibly oversee its deployment in their company. You also need to be educating the whole C-suite. This is not just a CTO or a CIO's business. As I say, you need to be taking the whole company with you on your AI journey so you can create those diverse teams.

**Katherine Kroll (21:18):**

You've said AI benefits from companies taking advantage of the attention economy. Is there a way that companies can still profit and help heal the attention economy?

**Aza Raskin (21:32):**

I think there are definitely places we can create a race to the top. Trust is going to be the most valued brand attribute in the future. You can already see Apple moving in that direction with all of their privacy ads. But just think about it. GPT-3 is just at the cusp of being able to generate text, which you can't tell whether it was written by a human or not.

**Katherine Kroll (21:58):**

GPT-3 is a new AI language model that can generate human-like text.

**Aza Raskin (22:04):**

Deepfakes are exactly at that level, too. That means we are going to be living in this, not uncanny valley, but synthetic valley, where we do not know what is real, what is not. Like the real world is becoming more virtual. I think there can be a race to the top for the companies that really own trust and demonstrate trustworthiness. Right? Like humane technology helps you love the way you spend your time and the decisions you make. I think everyone would agree that technology is not helping us make decisions that we love or spend time in a way that we value later. That, I think, is the opportunity.

**Aza Raskin (22:48):**

The challenge is that we are stuck in a rival risk game dynamic situation. Right? If Facebook or YouTube doesn't include auto play and TikTok does, it draws people in to TikTok. If TikTok doesn't figure out ways of generating that next perfect dopamine-releasing video, then some other platform will. As long as we don't have protections at the government or regulation level, that race to the top will never happen.

**Aza Raskin (23:27):**

As long as we have rivalrous game theoretical dynamics on top of a finite playing field, our minds the planet, with exponentially powerful technology, technology that gets more powerful over time, each step it moves up because it is using itself to improve itself, we will not gain a humane technological future. We will instead continue to extract and atomize and commodify the human experience, turning human behavior into dead slabs of predictable human behavior. That's why we need things outside of the market to bind it or to create guardrails, which honestly is not particularly revolutionary radical. Right? We've decided that, this is Shoshana Zuboff's phrase, but we've banned the sale of human organs and we've banned the sale of human orphans. It shouldn't be so surprising that we might also ban the sale of human behavior and intention.

**Katherine Kroll (24:33):**

One of the problems that feels a bit more unique to AI is that the stakeholders with the information necessary to regulate AI are the same people creating it. How do you think about that challenge?

**Aza Raskin (24:47):**

I think culture is upstream of politics, so I think movies like *The Social Dilemma* are really great levers for shifting popular awareness, because the tobacco industry I think is a really good example. Now, tobacco is another addictive technology. It doesn't get more addictive over time like technology, so it's not a perfect analogy. But there, a class action set of suits created a huge pot of funding for public service announcements, and that consistent and at-scale funding of culture or paradigmatic shift really changed the rates of smoking in the US.

**Aza Raskin (25:36):**

You could imagine something similar happening in technology. You could imagine something like in California, power companies, utilities, there was a perverse incentive to get you to use as much energy as possible. The solution was that, above a certain amount of energy use, your rate actually doubles. But all of that profit, the original and the double amount, gets reinvested into renewables, essentially into competitors. I think we could imagine a similar kind of attention tax where Facebook, above the amount of time that users say they don't regret using Facebook, which is—I can't remember the exact number—it's like 20 minutes. Something in that range. It's different for each platform. Any money they make from advertising, et cetera, above that 20 minutes, doesn't go to them. It goes into funding public digital infrastructure that has at the center the common good, or funding the fourth estate in journalism and the immune system of our democracy and creating an alternative. Because what's so inhumane is that we're forced to use systems that we need for connection, which is a fundamental part of the experience of being human, and that's inhumane. That's where we need governmental or regulatory protections to fund a new public digital infrastructure.

**Aza Raskin (26:58):**

This again is not particularly radical. My favorite example of this is Sesame Street, which most people forget was government-funded. It was an incredibly creative group of people coming to make public media. Imagine if we had that kind of public media, but for social media. When Sesame Street was created, there were groups of experts that came in on childhood developmental psychology, on pedagogy theory, on Muppets, and they came together to create something that was exceptional and beautiful and did bend the arc of history towards justice. I think we could do the exact same thing with our humane future.

**Katherine Kroll (27:40):**

Obviously, at Brown Advisory, we are investing in companies that we think are leaders on the issues that we've talked about today, but no company is perfect. From your perspective as an expert on AI, what is one thing you would urge investors to ask companies that they're invested in when it comes to ethical AI

PART 2 OF 4 ENDS [00:28:04]

**Arathi Sethumadhavan (28:06):**

That's such a cool question. No one has asked me that before. I'd say, the first thing is pay attention to what's happening, right? The forthcoming regulations in the AI space and all the discussions that are happening on this topic, because these certainly impact the companies that are developing AI solutions. You may be familiar with the recent EU commission's legal framework around AI, so that was in the news recently. So what that's going to do is it's going to set limits around the use of artificial intelligence in certain areas, like self-driving cars, like hiring decisions, scoring of exams. It may also even go to the extent of prohibiting the use of AI in certain areas.

**Arathi Sethumadhavan (28:50):**

So this is only the beginning, right? I would tell investors to pay attention to the regulatory landscape and articles and discussions in that front. So that's the number one thing. Number two, I would say pay attention

to what the companies that you are investing in are doing in this front, and how important is that sort of work, is ethical product development considered by the senior leadership team of that company? That's super important, because all of that correlates to consumers' trust in that brand and that might all have an impact on the overall performance of that particular company.

**Katherine Kroll (29:30):**

If you had one piece of advice for investors, would it be to ask companies how they're leveraging these AI tools? Or are there other things that we should be in communication with our holdings on?

**Kay Firth-Butterfield (29:46):**

Yeah, absolutely. So my first thing, if you don't ask any other question, you should ask, do you have an ethical or responsible AI policy for the use of AI in your company? That seems to me to be the most vital thing. And if they say no, then I suggest that you don't actually invest in them, if they're using AI. What does the ethical policy look like and how do they ensure everybody in the company knows what that policy is? How do they know that everybody in the company understands what AI is? And what are the implications of using AI to the company, to the employees and to society depending upon how you are deploying that AI.

**Kay Firth-Butterfield (30:41):**

I'd also suggest that you go on and say, "Okay, how are you ensuring that your product is ethically designed, ethically developed and ethically used?" Or you could just say, "How are you going to deal with the problems that are going to arise from your deployment, use and design of these AI tools?"

**Kay Firth-Butterfield (31:06):**

I think you might want to say, "Have you got a chief AI ethics officer or someone overlooking responsible innovation with AI?" Many companies have just got a chief AI officer. That person is usually a technologist, and what we know is that if you are going to be deploying AI wisely, then you actually need those diverse teams. You need somebody who can think beyond the technology for your social and governance needs around AI. I believe that it is the companies that take this seriously that will be the most successful as we move into the future.

**Aza Raskin (31:58):**

I think both as investors and shareholders, there's a lot of activist work that can be done. One is asking for social impact reports. These are not really well used these days, but just like there's an environmental impact report before you do some kind of major development that lets you think through beforehand what are the impacts going to be, we should be doing social impact reports. So as an investor, that's something you ask of a company. Have you game planned through what will happen societally by making these changes? By using this algorithm? Have you tested it?

**Aza Raskin (32:36):**

I think another one is asking for red teams. Just like in security, any company that wants to act responsibly should invest in red teams, groups that come in and intentionally try to understand where it'll go wrong. I mean, investors might not necessarily like this, but it creates a little bit of liability, because it means you've gone through, you've thought through it, you've documented it. But it also means you head off those negative externalities before you get there.

**Aza Raskin (33:01):**

And then the last one is just looking at disparate outcomes. Actually look to see how a product is affecting the world. Don't invest in the companies that are breaking the public commons. Business relies on a well-functioning society for profits, and if you take short-term gains at the expense of a functioning democracy,



well, you're not going to be making profits in five years from now, 10 years from now. So that slow or long-term view I think is the necessary one.

**Katherine Kroll (33:36):**

I think part of what AI represents is a super human power that we just can't compete with in terms of scale or scope or speed. And I mean that less in terms of misinformation or the attention economy, and actually looking at what a just economic transition will look like as we see a lot of jobs be automated away. I'm wondering if you've been spending time thinking about that, and if there are any examples of companies that are doing a good job?

**Aza Raskin (34:21):**

Job as a term is still fairly new. It's just a couple centuries old word. And I think the opportunity that AI affords us, and this is in the best possible world, is starting to ask the question, what is it to have a fulfilled and rich life that's independent of toiling or working at a job which can easily be replaced with software? And that's a fundamental reorganizing of society, and it's going to happen whether we want it to or not. Since the enlightenment, the prevailing paradigm has been the voter knows best. The customer is always right. Trust your heart, trust your feelings. And in a world of synthetic valley and deep fakes, and as you call it, this superhuman ability to make predictions about human behavior.

**Aza Raskin (35:22):**

When you sit down to play Garry Kasparov in chess, the reason why you lose, or let me just say, I lose, maybe you're amazing, is because he can see more moves ahead. He can out-compute me. And when Garry Kasparov sits down to play against a super computer, he loses because the computer is out-competing him. That's happening in every aspect of our lives. So we are losing the ability to say with confidence that the customer is always right, that the voter knows best, because we are living in a new era of persuasive technology. And I think that that's a very challenging problem to solve, because we don't know what's on the other side. And I know that's not directly answering your question of which companies are doing this well, but I actually don't know of any that have really tackled this problem full on, because they're still operating from the older set of paradigms.

**Katherine Kroll (36:22):**

In recognizing that there's no picture-perfect company that is knocking it out of the park, are there best practices that we can at least look for from companies? Are those intangible things that we can latch onto as we pave the gap between where we're at now and where we need to be?

**Aza Raskin (36:44):**

Great question. So here's some practical things that one can do. One, as a principle, people closest to the pain should be close to the people with the power. That is, the vulnerable or disproportionately affected populations should have a seat. There should be somebody who's a board observer, or on the board, who represents the most vulnerable, because that keeps you honest. Two is setting up defense and measurement infrastructure. Cities spend at least 25% of their budget on police. If you're making a social system, you should be expecting to spend at least that much on protections for the populations that you serve. Information companies, Facebook, Google, Twitter, are de facto governments, right? There are ambassadors now for most nations to Silicon Valley and that's not an accident. So we need to treat the companies as such.

**Aza Raskin (37:44):**

And so budget size, I think, is really important to look at when it comes to defense and protections. Thinking about the structure of your social system versus, say, moderation of your social system. So for example, if Facebook writes a classifier for identifying, say, COVID misinformation, well that classifier only works in

English, which means that the most vulnerable countries where their engineers do not speak that language, never see the benefits.

**Aza Raskin (38:20):**

I think an easier way of thinking about things, or instead of trying to write moderation tools, change the fundamental flow. There's some research that shows if you just limited the number of shares on Facebook or WhatsApp to two hops instead of infinite, that does more to diminish mis- and disinformation than anything they've tried so far. Of course it hurts virality.

**Aza Raskin (38:44):**

And then finally, think about thick data and anti-KPIs. That is, instead of just looking at metrics and reporting metrics, include qualitative data, real user research where you get into the field and you talk to users and communities, because metrics obscure so much. And two, anti-KPIs. For every KPI that you have in your company think about introducing an anti-KPI which keeps you honest. Instead of measuring just engagement, start to measure regret. At what point do users start regretting being on your platform? So pair every metric with an anti-metric so that you can't blindly follow and optimize.

**Katherine Kroll (39:29):**

Thank you, Aza, Arathi and Kay for sharing your knowledge with us and for the work that you do. I'll be back in a minute with two of my equity research colleagues to discuss our investment takeaways. Aza, Arathi and Kay shared some big ideas that I need the team's help to translate into what it means for us as investors. John Canning and Victoria Avara have been my go-to experts when unpacking this topic. John covers technology companies as a fundamental analyst, focusing on the financials, and Victoria focuses on the same companies to understand how environmental, social and governance factors impact those financials. As you'll hear, we agree with our guests that diverse perspectives and teams help us make better investment decisions.

**Katherine Kroll (40:21):**

Victoria and John, I am so happy you both are here. It feels very appropriate to be with you, because we've spent a lot of time together, huddled, trying to make sense of this issue, given your sector coverage as our fundamental and ESG tech analysts. And it feels like a very natural extension of the conversation we've been having for the last three or four years. So let's get into it. John, your voice was in my head during these conversations, both in agreement and in healthy pushback. What did you hear that resonated? What reinforced how you look at technology names from a fundamental lens? And also, what was missing from your perspective?

**John Canning (41:05):**

We heard from all of the experts in different ways about how these systems are black boxes, and it's very difficult to understand exactly why they're making the decisions they are. Yet at the same time, we're asking as a society for explainability and accountability of the systems. And the reality is these are probabilistic systems. Today artificial intelligence is really just advanced machine learning, which is really just advanced statistics, and it's not going to solve all our problems. And when we interact with our companies that we're invested with, or are evaluating new opportunities, particularly in technology, we're hearing about AI all the time. But I think what we really need to do is understand what they're trying to accomplish with AI and why. We need to move beyond the buzzword and understand, what is the customer problem they're solving? How are they doing it? And perhaps something that wasn't discussed a lot, is that, like any investment there are risks, but there are also big opportunities with artificial intelligence.

PART 3 OF 4 ENDS [00:42:04]

**John Canning (42:03):**

There's huge workforce efficiency gains that may happen over the next decade by automating repetitive tasks. We might be able to use artificial intelligence to do predictive maintenance on critical infrastructure assets that would allow better uptime and much freer movement. We can drive improving health care outcomes.

**Katherine Kroll (42:20):**

Victoria, what about you? From your seat as our ESG tech analyst, what landed and where would you add on?

**Victoria Avara (42:28):**

My brain is certainly buzzing with ideas yet I keep coming back to one of your questions to Kay. It was what would Kay say to someone who argue that investors should not worry about AI ethics? And Kay just said it comes down to one word, trust. Boiling this all down to trust feels so simple and so compelling. Then on the other hand, Aza explained that one of the major reasons why AI is mistrusted is the black box character of machine learning. John, we've been discussing this a lot in our investment discussions as the current level of transparency on algorithmic systems can limit our ability to assess the effectiveness of the governance practices. I'll say I'm encouraged by the fact that some companies are already finding ways to build trust in AI.

**John Canning (43:15):**

The importance of those teams to mitigate risks, Victoria, they can avoid really consequential outcomes that may or may not be intended. What are you looking for in the cultures of these teams as you do your due diligence? When we meet these companies, how are you assessing our comfortability with the controls that are in place?

**Victoria Avara (43:35):**

There's a lot that we could unpack here, but I think algorithmic bias is a good example of an AI risk that underscores why it's critical to assess corporate culture. Katherine, something I took away from your discussions is that a major way bias can creep into algorithms is through the human driven processes underlying the algorithm. So when we're looking to assess this and the human driven processes, as one part of our broader due diligence, there's a bunch of questions that we can ask. Does the company have of a responsible AI framework or processes that specifically aim to mitigate bias from design to execution? How is the company allocating capital and resources accordingly? Is there diverse representation at the table from design to execution? What accountability mechanisms are in place?

**Victoria Avara (44:20):**

Importantly, we're looking at who's accountable on the executive team, the board and throughout the company. And beyond that, how is the company monitoring their AI and machine learning models to detect and mitigate algorithmic bias after deployment. And John, that's something that goes back to your comment. I think red teams was something really interesting that Aza brought up. What did you think about that comment?

**John Canning (44:41):**

I really liked this point and we think about artificial intelligence amplifying content and the really negative impacts that can have if not controlled appropriately. But at the same time, it's not creating the content. It's not creating the issues that we have in society. And so I think when we think about these issues, we also think about where is the technology being pointed? What is the problem that's being solved? And what is the design of the AI? What is the design of this system that's being built on a massive set of data? We've talked a lot about the idea of the outcomes of these systems may have unintended impact, but they're not externalities. They're controllable at the end of the day. And it's up to the companies to ensure that they're comfortable with the outcomes that they're delivering to their customers and to society as a whole, as a result of these systems. So I love the idea of using a red team to really push the boundaries of what could go wrong.

**Victoria Avara (45:42):**

Yeah, I think that goes a lot to the diversity, equity, inclusion in our view on how we're thinking about that through culture and how that's really important to have someone with a diverse background on the design team or sitting there at the table, at least.

**Katherine Kroll (45:55):**

It's interesting. I see that piece about trust and teams that you dissect, Victoria, and that piece about unintended consequences and accountability that, John, you bring up really going hand in hand. Difficult conversations can lead often to better investment decisions. And I think that we're lucky that our ownership structure enables this type of culture that embraces challenging discussions because we're all owners in the firm, but it's still not easy to push back on a colleague about any issue, and certainly not when that issue is AI.

**Katherine Kroll (46:33):**

And then of course the very tempting urge to reach for blame instead of accountability, blame doesn't get us very far. Accountability can actually help us evolve and find those solutions that are so critical to us getting AI right. So I'm left wondering, John, what does it look like for us as investors to hold companies accountable for some of the unintended consequences that you outlined?

**John Canning (47:03):**

Yeah. At the end of the day, what I focus on probably most outside of some of the competitive barriers that a company might have is what is the problem that is being solved for the customer or the consumer of the technology. And when I sit here and think about holding people accountable for unintended consequences, I think we also should ensure that we're holding them accountable for the intended consequences of the systems they're designing and the problems that they're solving with this technology.

**Katherine Kroll (47:34):**

I'm reminded of something Aza said, and I have it pulled up so that I do it justice. He shares, "It's not AI or technology is an existential threat to humanity. It's that the worst of society is an existential threat to humanity, and that technology is amplifying that." So this challenge of AI amplification is a really big one and it has, and I'll use a not so technical phrase, scary consequences. That part is real, but it's also true that some companies are trying to manage those risks and that oversight can be really complicated. Can you help us grab it? Can you give us an example of an investment that you cover that's doing a good job at managing these risks.

**Victoria Avara (48:25):**

One company that sticks out to me is Microsoft. They've really built out teams and processes around this. They actually created a mandatory introduction to responsible AI course for all of its employees that includes a standard as well as building blocks of Microsoft's AI principles. So they're specifically embedding that throughout their entire culture and their employees something we're thinking about is when there are these risks that occur, how is the management team reacting, or what are they doing after that to make the controls and processes stronger. And that's what we really feel comfortable in the governance teams and the investments that we're thinking about is how are they reacting? Of course, there's going to be risk. Of course, AI is not going to be perfect, but how are companies continuing to make this better and stronger and mitigate the risk that we do see.

**John Canning (49:14):**

There's a lot of responsibility that lies with the biggest technology companies. But I think there's also a really big responsibility for us all to understand that at the end of the day, they can't control everything.

**Victoria Avara (49:27):**

I think that's a great point on the larger companies setting the stage for some of the smaller companies as well. One really interesting example was IBM launching an opensource toolkit that can help other companies examine report and mitigate discrimination and bias in machine learning models throughout the AI application life cycle. So how can these large tech companies provide and collaborate with others in the industry? I think that piece of it's going to be so important because at the end of the day, one company is not going to solve this.

**Katherine Kroll (49:57):**

So we've discussed how responsible AI demands that risk management. But I think the other side of responsible or ethical AI is investing in companies who are using it to solve for some of our world's greatest challenges. This is where we spend most of our time as a research team, finding compelling companies that are, let's call it, mission critical in their business model. And increasingly, we believe in order to be resilient, to be a compelling long-term investment with positive outcomes for shareholders, they need to participate in the solutions to these challenges. And we see that more and more being rewarded by the market. So when you think about our portfolios, where are we investing in companies at that intersection of strong fundamentals and sustainable or social opportunities that add to performance.

**John Canning (50:57):**

One company that we've been invested with a long time and we have a great deal of respect for as operators, but also in the manner in which they act in relation to society is Intuit. And for those who don't know, they have a variety of different businesses. The businesses span from small business accounting and payroll software to consumer tax-oriented solutions to now consumer finance with their recent acquisition of Credit Karma. And I think there's just a lot of really cool things that they're doing, leveraging artificial intelligence. One example is what they call the cash flow planner, where they're helping small businesses understand where they might have cash flow needs, and hopefully, improving the success rate of small business.

**John Canning (51:39):**

Today, 50% of small businesses go out of business in five years. And so if you think about leveraging the data that they have about businesses, their working capital needs, and the ability for them to forecast and help these small business owners who often don't have the resources to maybe build out a very complex financial model.

**John Canning (51:59):**

You talked about climate, Katherine, and for both of us, one of our favorite companies to talk about is Autodesk. They're very involved in most of the vertical assets or large office buildings that you see constructed across the globe really, but particularly in North America. And so what they're trying to do is to drive efficiency in the construction process using different machine learning opportunities. And if you think about 30% of construction materials going to waste during a project, there's potentially really exciting, positive implications on waste reduction, energy efficiency and potentially climate.

**John Canning (52:34):**

The last one that I'll bring up is a company that Microsoft is acquiring called Nuance Communications. They are prolific in helping improve productivity of physicians in a variety of different manners, leveraging conversational artificial intelligence. And so if you think about the need we have to provide care to an aging global population. And at least in much of the globe, Microsoft working with Nuance to bring productivity to the doctor, to help improve the number of patients they can see and hopefully the clinical outcomes as a result.

**Victoria Avara (53:08):**

Another thing I think about is that tech companies leveraging AI to create energy efficiency. So I think about Google's DeepMind division that's developed AI that teaches itself to minimize the use of energy to cool Google's data centers. And another area that I think about a lot is agriculture. So autonomous systems are already improving outcomes in key agricultural processes from planting to fertilization and pesticide use, and AI can collect and analyze data on variables such as soil conditions and crop health to boost the final output.

**John Canning (53:44):**

Internet of things is still so nascent, and the data that we're collecting outside of the digital systems that we interact with that are in the real world, that will be enabled by 5G and other technological advances is really going to explode. And as I alluded to earlier, at the end of the day, these are just probabilistic systems that need a lot of data. As the world digitizes, as you see companies move their infrastructure to the cloud, as you see internet of things explode with 5G and other capabilities that are going to drive autonomous systems, whether it's in the car or elsewhere, we're really just getting started on the data collection front, which means we're probably just getting started on the problems that we can solve leveraging that data.

**Katherine Kroll (54:30):**

I just feel very lucky that in many ways, this is another day at the office for me, where I get to listen and learn from both of you.

**Victoria Avara (54:38):**

Thank you both. That was great.

**John Canning (54:40):**

Thank you, Katherine. It was an awesome conversation.

**Ken Stuzin (54:45):**

Thank you again for joining us as we continue this effort to seek out insights that help us understand our rapidly evolving world. If you enjoyed listening, we encourage you to subscribe to the podcast. We are going to take a short break for the holidays and we look forward to being back here with you early next year. Until then, be well and stay safe. Also, we at Brown Advisory extend our best wishes to you for a joyous holiday season and a prosperous New Year.

PART 4 OF 4 ENDS [00:55:27]