



Beyond bias: Bringing behavioral science to hiring decisions

Host: Burt Rea, Managing Director, Deloitte Consulting LLP

Guests: Kate Glazebok, CEO and Co-Founder at Applied
Jim Guszcza, Chief Data Scientist, Deloitte Consulting LLP

Burt Rea (Burt): Talent acquisition. Recruiting has seen an explosion in technology tools with some 2000 vendors in the space. According to a large proportion of this year's global human capital trends respondents see the role of technology increasing across a range of recruiting processes over the next three years. Using data to find, source, and select candidates more efficiently and taking a data-driven approach to hiring is one of the recruiting function's biggest opportunities. Here to talk with us on today's episode of Capital H is Kate Glazebok, CEO and co-founder at Applied. My Deloitte colleague Jim Guszcza had the opportunity to talk with Kate about using technology to remove bias in hiring and improve hiring decision-making.

Jim Guszcza (Jim): Welcome everybody to our Capital H podcast. My name is Jim Guszcza and I will be talking with Kate Glazebok, the CEO and founder of Applied. Kate, welcome to Capital H.

Kate Glazebok (Kate): Thank you so much, Jim, really excited to be here.

Jim: So, let's start Kate with by telling us about you and your background in Applied and we'll pick it from there.

Kate: Sure, so I guess I came to Applied through an unusual path. Applied is the first technology spin-out of the behavioral insights team, just sometimes warmly and not too warmly referred to as

the nudge unit. So, we are the first technology to come out of it thinking about how do we apply the science of decision-making to hiring decisions and actually build an HR tech platform that puts quality of decisions first, not just speed. So, I think spending my career up until that point, focused on working with governments around the world and the UN on various aspects of social mobility, and I just basically stumbled into the very large research base, that many of the audience will be aware of around how quality of decisions in hiring. I guess nowhere near where we would want them to be and in fact displayed loads of these systemic biases, the behavioral scientists spent a lot of time exploring and in fact found that there was a huge area where we were making decisions that resulted in otherwise perfectly capable and in fact very talented people being overlooked for jobs just because they didn't quote, unquote "look the part" and I guess it just got underneath my skin and I couldn't shake it and so I started thinking about ways that we could use that science to turn hiring practice on its head and that's really how we came to build Applied.

Jim: It's fascinating, so in the sense of translational behavioral science, let's take what we have known from behavioral science the past four decades or so and actually create a tool, people used to apply in practice, is that a reasonable paraphrase?

Kate: Yeah, absolutely, in fact I think where we got to was we were really frustrated in saying the same kinds of studies come out time and time again sort of validating everything we already knew and there was this growing evidence base. And yet there was a huge, I guess, gulf between where the research science was that and where everyday practice of people making hiring decisions was and the real link was where they are not using technology that helps them to make the right kinds of decisions and that felt like the most important moment we could try and get involved with and trying, to use that behavioral lingo, to change the choice architecture of the people really face when they make hiring decisions.

Jim: That's fascinating. What we can do is, why don't we dive a little bit into the science, I think it's just a level, I think some of our listeners maybe familiar with some other science, there might be something that might surprise people. So why don't we spend a little time talking about some of these findings that you are talking about and then we will put it back to your software and how your software makes it easier to apply the science and makes smarter decisions in hiring and beyond. I believe that a lot of this goes back to work with Daniel Kahneman and Amos Tversky just recently popularized by Michael Lewis's book "The Undoing Project" and Kahneman's own book "Thinking, Fast and Slow" and I believe that distinguish between bias and noise, you want to talk a little bit about underlying research?

Kate: Yeah, absolutely, so I guess bias is one of these words that's now taken on this form, we started to see it now popular news outlets and I guess people are quite familiar with the aspects of bias like stereotype biases. So what happens in our brain when we see a person in a setting, we are not expecting to see them in, people might be familiar with the kinds of research around the fact that the brain takes ever so slightly longer to make the

connection between the word surgeon and woman than it does between the word surgeon and man just because we have absorbed the world around us, which contains many decades worth of skewed decisions that have resulted in us not seeing people in the same kinds of roles all over the economy. So, the kind of stereotype bias might be one of the things that people have come across in their everyday lives, but you are absolutely right. Daniel Kahneman and others have tried to help us distinguish between the concept of bias, which is to systematic and therefore to some extent predictable, once you understand the underlying makeup of it. And noise, which is in many ways random, so I guess to steal one of Kahneman's ways of describing this, he talks about a few to use an everyday sort of weights recorder. So if you wanted to make sure you are kind of reading a scale for your own weight on a daily basis, we might describe a scale has been biased, if it every day describes it has been 5 kilograms heavier or lighter than you actually are and it's predictable in that sense, we know that it is over- or underestimating your weight in a predictable way. But it might be considered to be noisy if on Monday, it tells you are 50 kilos, and on Tuesday tells you are 72 kilos, and on Wednesday it tells you are 65, in ways that wouldn't be possible. So, it's actually just unpredictable and noisy in its estimate of your particular weight. And I think this is hugely important because while it gets underneath the kind of wonky things we might think about in behavioral science, it's really important for decision-makers to distinguish whether the decisions that they want to improve or they want to improve them because they are currently noisy and the people are getting different outcomes than they should be getting and that's happening in a seemingly random way or whether that a systematically bias, so we systematically making decisions, which are over or underestimating risks for example over or underestimating temperature or over or underestimating in the case of what we care about the quality or perceptions of talent of other people. So, I guess Kahneman and his colleagues have done a lot to popularize these days of bias and noise, and also to help us distinguish different types of biases and how we might get underneath them.

Jim: And it's also important to emphasize, these are really two independent sources of errors, one is our minds are not like perfect computers, we don't have unlimited process and power, we can't handle unlimited lots of information, so we take a little bits of information and we tell stories with them and those are called heuristics. And those heuristics is the big a-ha discovery over the last 40 years, these heuristics can contain systematic biases and it's not just stereotype bias, a lot of biases like recency bias or the halo effect, I think the effects, I think if it's something you fundamentally like you tend to downplay the risks of it and vice versa definitely you don't like, you can play up the risk, there is no logical reason to do that. In stereotype bias, you read about in the headlines all the time, there's like dozens of these cognitive biases, but even once we've taken that into account, this still as independent source of error, which is noise, is that correct?

Kate: Absolutely, so to make that real I guess in the context of what we do, which is helping organizations to remove bias from the way that they assess candidates. We have actually run numerous

types of experiments to get underneath where does this bias come about. So, as you say the stereotype bias is also concept of infinity bias, which is that we all tend to get along better with people who are a bit like us and therefore we tend to like people who are a bit more like us. In a hiring context that can mean, we end up just building better rapport, which also has a sort of positive effect on the candidate's capacity to be at their best and perform at their best because there's already a positive rapport, so they appear to do better not only because we like them more and therefore kind of have a positive glow in the way that we interact with them, but actually they probably even do perform better. So, those are couple set of biases that we started to look at. We also started to explore other ways in which the decisions that we take might have other areas that come into it. So thinking about the context of that decision and whether or not we might say for example that you rate a candidate, exactly the same candidate, different maybe in the morning than you might rate it in the afternoon, which has the features of the noise that we talked about that Kahneman popularized and what we actually ran is a huge experiment testing how different people rate different candidates and whether the order that the candidate appears to you in your long stack of candidates to assess has any relationship to the score that you give them. So, what we want to do there was really unpack a very popular study that was taken, a number of years ago, on Israeli judges, that found that Israeli judges tend to be more harsh with their sentencing, the closer that they got to mealtimes, or another way thinking about that is that longer to be since they had had eaten and taken a break, they tended to get harsher and in fact sort of slightly more risk-averse in their decision-making and that resulted in them tending to be less likely to offer positive bail outcomes to defendants. And we wanted to know does this same thing hold true for candidates, which is to say like if a candidate can't control whether or not their application is reviewed at 9 a.m. or 5 p.m., does it make a difference anyway and we actually found three really astonishing things. The first thing we found is that most of us tend to be quite error-prone in our early sets of decisions, so we're really very noisy, we are calibrating an early stages of decision-making and that means that our scores are a bit all over the place and the second thing we found is in general, we tend to be more generous at the beginning of a review process and we become harsher with candidates later in the pile. So, equally good candidates just get a lowest score if they have reviewed 72nd than if they had been reviewed twelfth. Third thing we found which we thought was particularly astonishing, was that actually the scores that we give candidates reflect not just their own quality, but actually reflect how they compared to the last person that we saw. So, behavioral list of they call successive contrasting, which is to say what we've tried to call like the Nobel laureate effect. So, basically as a reviewer say a phenomenally good candidate, so let's say from my own organization that Daniel Kahneman would apply and he probably will get 100 score out of a 100. Whoever the unlucky candidate is that comes after Daniel Kahneman, I'm going to down-rate significantly just because no one could compare to him and actually he set a new high watermark in my mind and what we are going to...

Jim: Anchoring of that, right?

Kate: Exactly, it's a form of anchoring and referencing exactly and we found statistically you down-rate not just the next person, but the next three candidates as a result of that high scoring candidate. And the reverse is also true, which is that if you saw a poor candidate then the next candidate gets high score than they otherwise would. So, basically one of the things I take from this is, anyone who has a surname that starts with 'Z' should be very, very proud of themselves because they very likely to have always been reviewed last, which means they get the harshest kind of evaluations. Just what we discovered, there's tons of noise, candidates can't control how they reviewed and organizations would love to believe that their assessment of a candidate isn't down to random factors like what time of a day it is or the order the candidate appeared, so we tried to build that knowledge into the way that we build product. So we randomize all of these candidate applications later on sort of gets official.

Jim: This theme of noise you are talking about right now, this might be a real surprise for some of our audience and this is kind of independent source of error, I mean I have been studying as a self-taught person of behavioral economics for a long time, I think you are the first person to introduce this depiction of me. Kate, I remember I was speaking with you with the behavioral exchange at Harvard, 2016, we were having a coffee at Darwin's coffee shop and I remember I was speaking with you, those are real a-ha moment for me and I have been reading up on the this ever since, so I believe that Kahneman himself and Sunstein are collaborating right now on a book on noise, on this very topic and it's a hugely important and under-recognized source of error, in people's everyday decisions, they can very readily be corrected using algorithms, is that right?

Kate: Yeah, exactly, and I think what we need to start doing and, you know this better than anyone Jim, when we say algorithm what do people think? Most people go straight to a robot containing more AI than we actually know, actually exists in the real world and they tend to clam up at that very point. But actually, algorithms are really just ways of organizing mathematical information so that we arrive at a particular outcome and many algorithms don't need to be the kind of black box machine learning that a lot of people fear, we at Applied use algorithm that go simple as what I described earlier. So, we know that ordering effects are real, we know that this sort of review people differently based on when they appear, so we take a little pieces of candidate applications and we randomize the order in which they appear to different reviewers. That is very basic form of algorithm that accounts for bias and noise that we had identified in decision-making.

Jim: I think people often forget that there are algorithms that are not machine learning algorithms. There are algorithms that are just based on prior knowledge and kind of organizing a way to make decisions in the structured way.

Kate: Totally, and I think there's an important distinction that's worth stating here which is that algorithms of that sort have a benefit over some machine learning algorithms in the form of what we might think of a sort of usability or even face validity, so one of the things we care a lot about at Applied is that we build products that real people use. So, I could develop, taking a long time, but I think we could eventually get to the stage, we get to a perfect point of saying here is the perfect candidate for you, don't question it, I'm telling you they are best candidate for you. But at the end of the day if everybody looks at that, and goes that's just doesn't feel right to me, I feel like I need to be involved with the people who join my team, they will reject that particular kind of recommendation and in fact the likes of Kate Macey and colleagues at Wharton have done a lot of work on algorithm aversion. So, this is our generalized tendency to kind of reject recommendations made by computers that we personally as humans believe we're better judges of, and people decisions are classic example of that. We trust our own judgments when it comes to other people and oftentimes for good reason, but even if not for good reason, we tend to be more likely to stop using computers where we believe that they've made even small mistakes. And I think when we talk about the distinction between more traditional algorithms and the more sophisticated machine learning algorithms, we really run into this interesting problem of who is actually going to adopt these kinds of technologies and if they don't feel involved in those kinds of decisions, no matter how perfect the outcome might be, you actually make no impact on the world whatsoever. And actually being able to go back and audit how did we get to these particular conclusions? So, in the case of risk assessment, if I can go back and say well here are the six things that went into the model and the reason we have arrived at the outcome of saying they were too high of a risk or a great risk for the organization is, I can trace that back to the various variables we collected is helpful not in for organizations being able to go back and audit their own decisions and improve them, but also means you can go back to a potential customer, explain to them why the particular outcome was arrived at. So, analogy for us has been in the hiring space, we also have a tool that helps organizations improve the way that they describe jobs. So, a huge body of evidence suggests different people are attracted to jobs on the basis of the words that you use. And so we have this concept that there is gendered language, so if I use stereotypically masculine-coded language or various stereotypically feminine-coded language, I can actually change the rate at which women and men applied to ostensibly the same job just by tweaking the kind of descriptors that I use for that role. So, if I use very agentic language, individual driven challenging role that will be more likely to have more men apply on average than if I use more communal language, saying this is a collaborative role we can have to work with a team to achieve certain types of outcomes, that's more likely to appeal to more women on average. And one of the things that we developed as part of that tool is that users of that tool can literally go and have a look at where the scores have come from, so they can see that they are getting uprated or downrated on the quality of their job description based on things we know researchers pointed to. And I think that has some benefits over other tools which might be more

machine learning led in that same space because people if they can't understand how you arrived at an outcome are much less likely the pick up a particular tool, and at the end of day we want people using tools we think improve its decisions, we have to pay some attention to this generalized tendency we know to want to be able to understand how we arrived at a conclusion.

Jim: Yeah, at a high level what you are saying resonates with me for another reason. There is a lot of talk in sort of the popular press, so people who are not data scientists, more like data science pundits, they always talk about "sophisticated machine learning algorithms." What you are doing is truly sophisticated because you are using algorithms in a very principled way to both to kind of like short circuit the biases that affect our decisions and the noise and you just described using behavioral design even to affect the pipeline like if the gendered language that prompts fewer woman to apply, let's change the gendered language, so use this behavioral design to actually change these preconscious decisions that applicants are making. But the use of "machine learning models" is very often not so sophisticated, if you just take a lot of data from historical decisions, train a machine learning model on it without thinking about the sampling train, you just kind of like encode and amplify those prior biases that's not very sophisticated at all, is it?

Kate: No, I think unfortunately that's the other big looming challenge that we all . . . in society face even though kind of rapid introduction of machine learning and AI into loads of decisions. I'm not a user of machine learning in our current product for reasons that we worry a lot about this kind of risk of algorithmic bias. So, to take it back into a hiring context, if at the extreme we have only ever hired people of a particular group and a particular background into a set of roles in an organization, using existing organizational data to help diversify that pool will fail because literally that data set doesn't contain any of those counter-normative groups on which to build and understanding of whether these groups might be a good fit for the organization and because we are worried so much about that and we were working with organizations to helping them diversify, they are hiring outside of those they previously hired that seems, I guess, counterproductive to say the least. But that being said, I'm actually a believer that there are some great auditing tools that are developing in the market that will help those who use machine learning to avoid the pitfalls of the sort of things we described before like what's the providence of that data, who has defined the questions that data contains, what's the missing data that we don't have access to, and how do we account for that kind of missing data. And then of course further on down the line, you can use tools like Themis developed by the researchers at University of Massachusetts-Amherst, but there are others that have come onto the market both in open source and paid-for form where you can have them almost meta-audit your own algorithm(?) to check for potential risks of bias in the outcomes that it derives.

Jim: Kate, let's pivot back to your tool, I mean I think we have covered the sort of the naive use of algorithms, the machine learning based algorithms can lead to unintended consequences and maybe artificial stupidity instead of artificial intelligence sometimes. So your

use of algorithms, it seems to be a nice way of doing it an end run a lot of these problems that are hitting the headlines. So, let's talk about how your software is designed to counter these two sources of error, bias and noise. Can you say few words about that?

Kate: Sure, so I guess our mission as an organization is to help organization's find the best people for every job regardless of their background. And it turns out, when you start with that mission, you really actually have to disentangle quite a lot of what is common practicing, most hiring practices, in two directions. One is a lot of the things that we test candidates on at a moment are not actually great predictors of their performance in role, so we tend to be using as you talked earlier about heuristics. So we tend to be using heuristics or proxies for talent, like where someone went to university or the previous job they were in, or the use of their experience, as proxies for whether they are going to be a great person for the job. And in fact, it's a very little correlation between those things and actual performance in the role. The first thing is helping organizations get at the skills that are really most likely to be predictors of those performance and actually testing candidates on the real things that they are going to do in the day job, which is the most predicted form of hiring. The second thing that we wanted to tackle is removing the noise and bias from the way in which we assess candidates. So, I guess to use the sort of tried and tested example of this is economists and psychologists have been running these experiments where they send out exactly the same qualified CV to thousands of employers by changing the name on the top of that CV and measuring the extent to which the name of the candidate predicts their likelihood of getting selected. And unfortunately the latest study just came out here in UK quite a few weeks ago that showed that people with nonwhite-sounding surnames in UK, for example, need to send up to twice as many CVs to get the same rate of contacts or interviews. They were equally qualified for the job, so not only is that depressing in and out itself, but what's more depressing is that they have been able to backtrack. The first version of that study was done about 50 years ago. Equivalent studies in the US have found that over about the last 40 years, there has been some modest potential directional reduction in the rate of discrimination against Latino and Latina groups, but no change in terms of African Americans. So, let's take that as a basic example, clearly there is something going on in our ability to assess talent when we get distracted by the names and signifiers, ethnic and otherwise, of candidates, so a really easy thing that we do within the platform is we get candidates answer work-relevant questions as the form of their application and then we put those answers through four types of devising. The first is we anonymize all of their applications, so people can't be distracted by knowing your name is Jim and my name is Kate. Second thing that we do is we actually chunk up the applications by question rather than by candidate, so we almost do a horizontal chunking process rather than a vertical chunking process. So, when candidates are being assessed, they have been compared directly to other candidates on each of their questions and reviewers stay just looking at well, who is the best person for the job when I assess their technical capability? Then I'm going to assess their ability to kind of work within teams to deliver a project on time and then I'm going to assess, let's say their ability to write a really great brief for a senior decision-maker. So,

that kind of horizontal comparison results in much faster decisions on the brain, because you stay on one task on one time, but also we can account for this halo effect, so if you do well in one area, no one is going to know that and they can fairly assess the other aspect of your application. The third thing we do is we, much like the Moneyball thing, we allow multiple people to score independently, they don't see each other's scores and then we just average them out because our data suggest, much like Google Zone when Laszlo Bock, when was head of People, they did the same study for Google where they found actually the average score of multiple people is a much better predictor than any single person's score would be.

Jim: Wisdom of the crowd.

Kate: Yeah, exactly. It's the wisdom of the crowd. So, no matter how expert you are, generally speaking, we'll get a better outcome if we ask you and a couple of other people for their opinion. And then the fourth and final thing that we do is we randomize for reasons that we described earlier which is that we found very strong ordering effects. So, if Jim you'd applied to a job, your answer to question #1 might be reviewed 1st, your answer to question #2 might be reviewed 22nd, your answer to question #3 might be reviewed 154th by different types of reviewers at different times. So, we take all of this data and we compile it back up and we develop a really, really rich data set on the quality of candidates in different areas of their application. Now what I've described might sound laborious, but actually it's no more information you're currently reading when you're assessing candidates. It's literally as much data as you get from a CV or a Cover Letter but it's significantly more predictive, and it's significantly less likely to result in a biased outcome. And what our data have suggested when we've run validation studies is that actually over half up to half or more of the candidates that you actually hire out of a process done through Applied you wouldn't have hired, if you've been looking at them through traditional means and that group is much more likely to be from a diverse background. So that's sort of some of the way – I guess remodel hiring processes around the real world of how people actually make decisions.

Jim: So, it's no harder to execute and it avoids the kinds of algorithmic bias you read about in the headlines, but it's a kind of algorithmic decision-making, and it ameliorates both bias and noise, and I think the four techniques you just described, it sounds like the first two sort of attack the problem of bias right like taking away the means that may bias you, or the kind of horizontal screening risk versus the halo of that effect of best practice. In the second two, the wisdom of crowd thing and the kind ordering those can tackle the noise issue.

Kate: Yeah, that's absolutely right. We've always been really keen that we tackle both because both have huge impacts on the risk of making hiring mistakes. At end of the day, most organizations just want to know they are finding the best people for the job and if most people are like us and they believe that talent does not know gender, does not know race, does not know background, then actually that should result in a far more diverse set of people for selected for jobs and that is in fact what we find.

Jim: Yeah, and going back to the Moneyball kind of analogy, it's barely been an analogy to them, Moneyball really literally a hiring story about hiring and using data to make hiring decisions. When you talk about gender and race, there's absolutely an important, kind of moral and ethical component to this, we want more diverse workforce that good for the kind of society we want to live in. So, it's a huge economic intention do this as well although it seems a lot of leaders might not appreciate that at this stage

Kate: I think that's right, and I would sort of put that down to two things. One is that it's fair to say that the evidence based on really why diverse teams tend to perform better has ramped up in the last 10 to 15 years. So, from about the last 20 years, we have known that there's correlational evidence that more diverse teams tend to be associated with teams that have better business outcomes. So, many people will be familiar with the work of McKinsey for example that have shown that more diverse leadership teams are more likely to outperform industry averages. So, if you have a more diverse by gender, you are more likely to outperform industry averages by around 15%. If you have a more diverse by ethnic background, it's up to 35% more likely. So, this sort of correlational evidence has been around actually for a little while, but we are now getting much, much closer to understanding what is the actual causal mechanism then. So, we now know through the work of people like Catherine Phillips at Columbia and others that actually diverse teams drive those better outcomes, not the other way around. And the reason why they seem to find diverse teams generate more creative and more innovative outcomes particularly in settings where information is unclear. So, the reason why diversity pays particular dividends is that by definition you have people approaching problems from different perspectives. And so, you are much less likely to run into the risk of group think and having everybody so just take your example of risk as you say, you don't want everyone who perceives risk in exactly the same way making that decision, one because it is redundant. If Jim and everybody are alike, just get Jim to make the decision. So, actually, we want someone saying, "Well, hi actually I think that this person might represent more or less of the risk because of this and from my background and my experience, I have a slightly different stance on the data that we are looking at here. So, some examples of studies that they've done has found that actually let's take jury decisions. They found that more diverse juries tend to arrive at more accurate decisions because they are more likely to remember clearly the facts of the case. So, they found that actually diverse juries tend to "work harder," their brains are working a little bit harder partly because we know we can't assume that everybody is making the same assumption and so we are almost bringing more of our own expert judgment to the table. And we are having to trash it out, we are having to say yeah, but actually you remember that we also learnt that the witness on Tuesday morning, arrived there at 10 o'clock, "Oh, actually I've totally forgotten." So, diversity and the evidence based on diversity teams performing better has become stronger and stronger and I think that's why we are starting to see particularly leadership teams pay more attention to the business case with diversity not just the moral case, which you and I may agree with, but at the end of the day, we also know a lot of people make decisions with cold, hard business cases. But I think to

your point about why we are not seeing this having happened. I mean frankly status quo is a powerful thing, we all stick to default actually deeply behavioral why we haven't seen the changes we have, but there is a huge gap between knowing there is a problem and actually being able to solve it when you are faced with the 70-second decision you've taken on that given day and it's 6:55 PM on a Friday night when you are trying to leave the office. And so really actually deploying behavioral design in the tools we use is I think the big change that should hopefully actually move the dial on these kinds of issues and that's really what we believe we are part of and we're looking forward to welcome lots of other organizations that are helping to embed this practice into everyday practice.

Jim: I am so inspired by that. And one last question, Kate, before we close and thank you so much for this wonderful conversation. Let's just leap ahead. Let's just say like two years from now, the majority of the Fortune 500 are using your software. That's what I want to happen. What is the future of this approach? I mean it seems to me this is not only for hiring, but you could also use that in principle right for promotion decision, talent management. You talked about senior leadership, you want more diversity in most senior leadership. Well, senior leadership often gets promoted from within. It seems to me that the same biases and noise probably can compromise hiring decisions, it can also compromise promotion decisions, right? What about using it for university admissions? To me there are a lot of areas where this approach could be applied.

Kate: Yeah, absolutely, and in fact we have spoken with a few universities about doing that and we are increasingly talking to organizations about I guess exactly as you say, promotion decisions. So, we've had over 95,000 candidates be assessed through our platform now. So, we know a lot more than we ever did. So, I think one obvious thing is we are constantly iterating and experimenting with our own data and how to optimize the process we developed. We are not going to stop with the four things that we designed here, and in fact we went from designing assisting process to also to building out an interview module because at the end of the day we want to make sure we are just not punting the problem one stage further down the line. We now have end-to-end solution that looks at everything from the way you advertise the job description and even down to the level of how do you deliver feedback to candidates if they don't successfully get the job, but in a way we know that they will learn from it and we are running an interesting experiment with Adam Grant at Wharton and his team on how to deliver better candidate feedback. So, I think there is a phenomenal amount that we can do with just better-quality data that delivers considerably more value just in the hiring space. But you are absolutely right. One of the things we've been starting to talk to organizations about is that they start to identify big gaps that appear in access to what we might loosely refer to a sort of the special project or stretch project. Here in the UK, we now have regulations that are actually organizations reporting on this very month to report on the gender pay gaps and they have identified that some of the areas where gender pay gaps have emerged is that maybe men

just incrementally have had a few more opportunities to be seen by senior managers and we all know that reputation is hugely important in most organizations and they've had a few more opportunities to be seen in those types of settings that are more likely to mean that they can consistently more likely to be chosen and promoted later on in their career. So, people are even starting about to use our product to help select people for special projects within an organization. So, it doesn't even have to be external hires, but even helping with an internal selection process. I think the university admissions problems is a very real one and we are actually excited to be talking to some organizations about that as well. Another area where people have come to us for is grant-making decisions. So, not only investors thinking about who they invest in which is equally an information problem, they have hundreds of companies come to them on a monthly basis asking them to invest in their businesses and in theory they should have a fairly structured set of questions that they are asking these businesses that might be amenable to over a structured assessment process, but likewise even grant-making bodies who deliver phenomenal amounts of money in terms of philanthropic grants and others, we think actually we'd like to be making sure that we are giving our grants to the widest possible group of people. So, I think in many ways you are right. What we've tapped into here is really about making assessments of other people and selection decisions. There's an obvious use case in the context of hiring because we make millions of these decisions every day and we have huge amounts of evidence to suggest we are not doing it particularly well, but I think the future in some sense is open for using better behavioral design through the many people decisions across the pipeline.

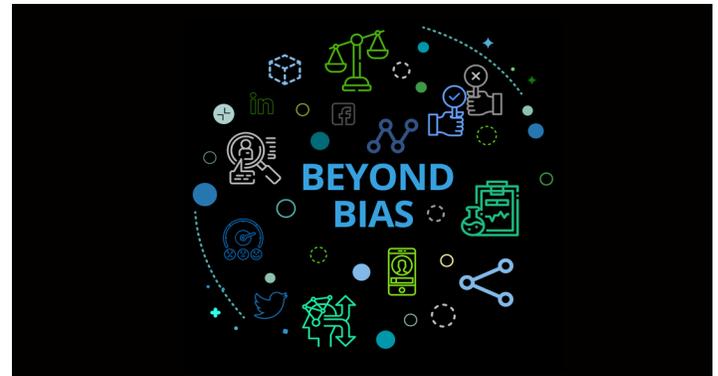
Jim: Thank you so much.

Kate: No, thank you very much for the opportunity. I mean it's always great tuning with you, Jim.

Jim: That's always a pleasure. That is why I have learned not just one, but I learn many things every time I talk to you, Kate. So, thanks again

and thanks from our audience as well. So, thank you again for tuning in and for future episodes of Capital H.

Burt: We'd like to thank our guest, Kate Glazebok, for walking us through how a client is using algorithms in new and different ways to help eliminate bias and cut through the noise in hiring to get at the right candidate for the job and to enable organizations to build a diverse workforce. Join us next time as we dive into more topics and trends that focus on putting humans at the center of work.



Learn more



Visit the Capital H podcast library
deloitte.com/us/capital-h-podcast



Explore our blog for additional insights
capitalhblog.deloitte.com



Join the conversation on Twitter
[@DeloitteHC](https://twitter.com/DeloitteHC)
[@Bersin](https://twitter.com/Bersin)

This podcast contains general information only and Deloitte is not, by means of this podcast, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This podcast is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this podcast.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.