The Future of Work podcast is a weekly show where Jacob has in-depth conversations with senior level executives, business leaders, and bestselling authors around the world on the future of work and the future in general. Topics cover everything from AI and automation to the gig economy to big data to the future of learning and everything in between. Each episode explores a new topic and features a special guest.

You can listen to past episodes at www.TheFutureOrganization.com/future-work-podcast/. To learn more about Jacob and the work he is doing please visit www.TheFutureOrganization.com. You can also subscribe to Jacob's YouTube channel, follow him on Twitter, or visit him on Facebook.

Jacob Morgan: Hello everyone! Welcome to another episode of the Future of Work Podcast. My guest today is Cathy O'Neil. She is a mathematician, data scientist, and author of a new book called the Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Today, we're going to be talking about her book and why relying or just pursuing data with a blind faith mentality can sometimes actually cause more harm than good. Cathy, thank you for joining me.

Cathy O'Neil: It's my honor to be here.

Jacob Morgan: Let's start off just learning a little bit about you before we even talk about your book. You have blue hair. I'm definitely going to ask you about that. How did you get into this world of math and data science, and writing this book?

Cathy O'Neil: Well, I've always been a math nerd ever since I was basically five, and learned about prime numbers. I've always loved the beauty of mathematics and the simplicity, and in fact, the cleanliness if you will, the way that Math is true even if you go to a different country or a different galaxy for that matter. It's not up for debate. It's just you have axioms, and then you have logical conclusions.

Jacob Morgan: I can't disagree with you that one plus one is two.

Cathy O'Neil: Exactly, unless we have a different definition of one or plus or equals. That was nice, especially in the context of learning about manifest destiny in third grade or something along those lines. There was only one right answer, and that was nice. I loved math for its beauty. It's almost an art to me. I wanted to be that kind of mathematical mind that sees beauty that no one else has seen. I became a math professor because that's what I wanted to do. I was a Barnard College professor here in New York City. Once I got here, I realized that I was interested in the real world too and being a business woman. That was in some sense far too organized and articulate and outspoken to live within a math department.

I got this first job that I knew how to get, and this is in 2006 I should say, which was as a hedge fund quant. I really had no idea what I was getting myself into.

Jacob Morgan: What is that for people that maybe don't know what a hedge fund quant is?

Cathy O'Neil: It's a quantitative analyst. It really depends on where you work, what you do as a quant, but I happened to work at one of the most quantitative hedge funds around called D.E. Shaw. I actually even worked with Larry Summers on a couple of his projects there. In some places like banks if you're a quant, you're working for a trader. On a trading desk, you're balancing the different kinds of hedges at the end of the trading day, which is a kind of low level grunt work type of scientific quant work, but in the company I was working at, the quants rule, and we have traders who do what we tell them to do or what our algorithms tell the machine to do really. Everything was all about being the smartest person and figuring out how to predict the future's market in my case.

Jacob Morgan: Are you a fan of the show Billions by any chance?

Cathy O'Neil: I don't know that show. I don't have a television, I'm sorry.

Jacob Morgan: Okay. We can get in all these fun facts about you too. There is a show on Billions, and there is a character there. The character is this math savant that works for this billionaire guy, and it's all about numbers and data and figuring out what to invest in, what not to invest in. I imagine you in that position of just tons of numbers and tons of data, and figuring out where to put money, and where to take out money at the right time. That's like what this is.

Cathy O'Neil: Yes, that sounds vaguely analogous. I mean, it was almost entirely statistics and numbers. There was very little, to be clear, very little investigation as to the business practices of a given company, whether we were going to invest in them or something along those lines partly because we were in futures, which is a slightly different market, but also just because as a quantitative firm, you're not talking about the qualitative aspects of the market. You're really trying to find patters in the past that you think will be continuing in the future that you can trade on.

Jacob Morgan: Got it. Where are you spending all your time now? What is a day in the life of Cathy look like? Why don't you have a television, and why do you have blue hair?

Cathy O'Neil: Well, I have to continue my story, which is that the financial crisis erupted, and I was quickly dissolution with what I was doing. In particular, I was very unimpressed with the way that mathematics was being weaponized, not in my firm directly, but in finance in general, the way that those mortgage-backed securities were being slapped with triple A ratings that we're absolutely not deserved. Triple A ratings were meant to suggest that rigorous mathematical analysis had been done, and these mortgage-backed securities were found to be extremely safe investments, but none of that was true. It was the first time I saw how much trust people had in mathematics, and how that trust could be absolutely abused.

I left finance. I started my blog, which I called Math Babe. The goal of my blog was to blow up what I considered corruption in the world of hedge funds and banks and finance in general. I needed a job, and I was still a nerd, so I ended up becoming a data scientist. I was in this startup world of New York City and ad tech. It was fascinating to me, because it on the one hand seemed almost exactly like the same job, because I was doing the same thing, predicting. Instead of predicting markets, I was predicting clicks.

Jacob Morgan: You were doing this for who, I'm sorry?

Cathy O'Neil: I was doing this for an ad tech company called Intent Media.

Jacob Morgan: Oh, for an ad tech company, that's right.

Cathy O'Neil: Yeah, Intent Media. I was trying to predict who was going to buy a hotel room on Expedia for example.

Jacob Morgan: Very cool.

Cathy O'Neil: But it really wasn't that different is my point. Moreover, the same kind of arrogance was there. Like, "Oh, we're doing this amazing work, and we're the smartest people in the world." Then I was like, "Well, what else could be similar? Like, could we also be weaponizing algorithms in the context of data science as we did in finance?" As soon as I started questioning that, I realized it was of course the case. The first sign of this weaponized mathematics outside of finance came to me through actually my college roommate who was a principal of Brooklyn High School in Downtown Brooklyn. She told me that her teachers were being scored with a fancy but secret algorithm that she couldn't understand.

Their tenure was on the line. She had a bunch of new teachers, so they had to get good scores so they weren't going to get tenure. I was like, "Well, this seems like a pretty important scoring system. Like, go find out what the formula is, and I'll explain it to you since I'm a mathematician." She asked for an explanation, but she was told by her contact to the Department of Education that it was math, and she wouldn't understand it.

Jacob Morgan: Nice.

Cathy O'Neil: Yes, I was like, "That sounds familiar, and not in a good way." That's the definition of a weaponized algorithm. Like, "Sorry, you're not ..." Similar to triple A ratings and mortgages. Like, "Trust us. We did the math. You could trust us." I was like, "Well, I don't trust anyone anymore." I started to look into that system that was assessing teachers. It wasn't just in New York City. It was in Chicago. The Chicago Teacher Strike was largely about this system. It's in Los Angeles, in Houston, and Washington D.C. The more I looked into it, I was like, "This is terrible, but it's a secret." It's being used to fire people in D.C. I found someone who got fired because of a bad score. There's no statistical evidence that it's meaningful, that it works, that it's fair.

Yet, it's being used in these high stakes' decisions. It was bad. At the same time, I didn't yet feel like I was participating in this. I wanted to be like a whistleblower in that sense. Until one day, I was working at this tech company, and a VC came to visit us, a venture capitalist who was possibly interested in investing in my company. VCs are powerful. They're powerful people. They get to decide which companies continue, which startups get to keep going, and which startups fold. I considered them like the architects of the modern internet in a certain way, because they are the ones who funded those innovations.

This particular guy, and I don't think he was unusual. It's just that he taught me something that I didn't like, which is that he told us that his vision of the future of tailored advertising, which is of course what I was working on, tailored advertising, his vision was that soon enough, he will see only trips to Aruba and jet ski ads, and he would never again have to see another University of Phoenix ad, because University of Phoenix ads were not for people like him. When he said that, a bunch of people laughed. I was like, "What just happened?" What does that mean? What does that vision imply? What I realized that it implies is that we have come a long way since we thought of the internet as a democratizing force, where information is free and everyone is equal.

That in fact, we were building a new kind of infrastructure, a new kind of environment within the internet, wherein people were segregated by class, by race, by opportunity. Then I was far from being a benign influence. I was contributing very directly to this technology that I was helping people figure out how to separate the winners from the losers. I was doing that by inferring, "Oh, what kind of computer are you using? What kind of browser are you using? What kind of purchase have you made in the past?" I realized this isn't just me either. This is everyone. This is what the internet does. It separates people into basically marketing silos. It propels the successful winners up. It pushes the unlucky people down.

I noticed all this, and then I also realized that I was going to be in the first category. As a data scientist, highly educated, PhD, white woman in New York City with extra money in my pocket, I was not going to suffer from these mechanisms, but I was building a system where it was quite likely that people were suffering, but that I would never see them. That's the other thing. It was invisible. That was really what was different ultimately from the financial crisis to what I now think of as the data science crisis that has already started, which is that in the financial crisis, we all noticed, because it was a huge mess, because Lehman fell, and then the markets planished, and everybody was worried about the functioning of the financial system.

Nobody didn't notice that unless they were actually not paying attention at all, but nowadays, if a data science algorithm or an AI, which is the same thing, it's just a new marketing ... If an AI program offers a job to someone who doesn't deserve it, or even worse, if it doesn't offer a job to someone who does deserve it, no one notices, not even the person who lost that opportunity that they deserved. In other words, the failures of the data science economy, the algorithms are almost impossible to trace, and so rather than assuming they don't exist, we should assume that they do exist, but we need to start measuring them.

Jacob Morgan: I have a million questions for you about what you just said.

Cathy O'Neil: That was a long thing, I'm sorry. I should have stopped.

Jacob Morgan: No, no, that was fascinating, but I know people are going to say, "Well, what about the no TV and the blue hair?" First, those two fund things, and then I got to dive into some of these things that you just talked about, because I have a lot of questions for you.

Cathy O'Neil: Why no TV? Simply because I don't think cable is a good way to spend my money. It's just too expensive. I subscribe to mlb.tv, because I'm a fan of baseball, and we watch a local television, because we have an actual antenna, an old-fashioned digital antenna. I mean, it's not that old-fashioned, because it's digital, but you know what I mean. We get free TV on the airways. We watch baseball on the computer, and obviously, I have Netflix. Don't get me wrong. I'm not crazy. In terms of the blue hair, I think, the easiest way to explain that is that as a 45-year-old woman with graying hair, I am essentially invisible to young people, but if I have blue hair or till hair, which is what my hair color is right now. It's actually called atomic turquoise manic panic.

With turquoise hair, young people talk to me. I'm visible to them. I'm interesting looking enough to break in to their perception. The underlying answer, the underlying truth is I actually really like young people, and I like talking to them. I just like also not being taken too seriously. That's the other great thing about having ridiculous hair. By the way, I don't think it looks good. Just to be clear, I think, it looks ridiculous, but that's what I am going for. I'm going for ridiculous because you are a bit ridiculous, and you are aware of it, and other people are too, then the stakes are lower, and you can have really fun conversations. I like having fun.

Jacob Morgan: I love it. I love it. I think, to unpack some of the things that you explored, maybe we should start with the very basics. That is what is Big Data? For people that are maybe not familiar with what's going on around this, or even if they are, I think, maybe giving an overview of this new world of Big Data would be really great. Then we can dive into some of those other themes that you looked at. Let's start with what is Big Data? Is this a new thing that's over the last few years? Haven't we always had data?

Cathy O'Neil: It's a good question, and it's a multi-part answer. I'll say this, big data is a term that is a technical term, but it's also a marketing term. I'm going to talk about the technical term first, but the marketing term is more interesting. The technical term is it depends directly on the technology you're using, but it used to be more data than you can fit on your computer. Now, it's more data than you can analyze in a few minutes. You need the Cloud or something like that. It just means you can't do a simple straightforward analysis, because it's just too much. It's probably also not in a database in a clean way, so you can't just perform a SQL command on it. You have to munch it in some way. You have to clean it. You have to put it into a form that it can later be analyzed.

That's a technical concept of big data. It's pretty vague, but, I think, that's what people mean. Then the marketing term is the term that I really focus on and object to, because it's truly misleading. There's a couple things that are misleading, and a couple things that are aren't. The thing that isn't misleading is that big data represents this concept, which is new. Yes, we've always had data, but we used to have to go around collecting the data that we wanted. The new thing about big data is the promise of big data is that you can collect data for one thing, and then reuse it on another thing. I know that sounds vague, but let me just explain. You can collect consumption information like browsing history and purchases and even possibly in real concrete stores, brick and mortar stores as well as online stores, you can find out who someone is as a consumer.

Then you can infer who they are as a voter or who they are in terms of their health. Are they living a healthy life? Are they at risk for diabetes? These two examples, inferring political leanings and inferring health risks from what looks like irrelevant information, consumption like purchases, those ar very, very meaningful connections. There's no obvious reason why that would work. Why should the way I buy things or where I buy things, and how often I buy things imply how I'm going to vote? It's not obvious. It's not a logical thought experiment that you can run that would make that clear, but it just happens to be the case.

Those are the two things that I worry about the most. If you wanted me to tell you, I worry about our democracy being threatened essentially by the fact that our browsing, our profiles as consumers are so well known, that our profiles as political thinkers, as informed citizens are vulnerable to propaganda, to fake news, to political messages. I also very much worry that, again, because our marketing profiles are so well known, there's very little protection on any of that stuff in the United States, that people will be able to infer how much we are going to cost in the future insurance costs. That will possible prevent us from getting hired at a job we really desperately want to or desperately need, because we present as high risks, which is expensive.

Those two concerns, nobody's ever talked me out of them, because that's what big data really I. Big data is a technology that allows us to collect seemingly innocuous data on the one hand, and then turn around and use it against us. Now, of course, proponents of big data wouldn't say it like that. They would say, "Oh, well, we get insights on people in other ways." Some of that of course is quite useful to us. For that matter, I do not object to myself being known for my consuming behavior. I really do get exactly what I want when I want to buy something on the internet. It's amazing. By the way, did I mention that I'm also very overweight? I love a blue hair, overweight, middle-aged lady.

When I was pregnant the first time, because I have three kids. One of them is 17, and one of them is eight, so I was pregnant in 2000. It was really hard to find maternity wear, because I would go to a store that has larger sizes, and they wouldn't have maternity wear. I'd go to a store that has maternity clothes, and they wouldn't have larger sizes. Then I got pregnant the third time in 2007, and dude, it was so easy to find maternity clothes. Do you see what I mean? The internet is a wonderful place for weirdos like me, people that are little bit strange and can't find what they need in the mall. It's like a nirvana for consumption. It is not a nirvana for other things. That's the problem and the promise of big data. It is that you can be served as a consumer, but be attacked and threatened as a citizen.

Jacob Morgan: Yes. I mean, you actually talked about a story in your book. I was hoping maybe you can share that, because I think, it really illustrates this point pretty well. Correct me if I'm wrong. It was in Washington D.C. It was the teacher's story, and I think the teacher's name was Susan, right? Then you talked about the algorithm that caused her to get let go. Do you know which part I'm talking about?

Cathy O'Neil: Sarah Wysocki, yes.

Jacob Morgan: Oh Sarah, not Susan.

Cathy O'Neil: Yup.

Jacob Morgan: Can you maybe spend a minute just explaining that story, because I think that it will really help people understand the dangers of purely just relying on an algorithm as far as how to make a decision.

Cathy O'Neil: I'll back up a little bit, and say what the scoring system was actually purported to do. There's a long tradition in our country of holding teachers accountable, because there's a bunch of presidents in a row, and Trump isn't one of them actually, who wanted to be known as the president that fixed education. That meant a lot of things to different people, but a consistent theme was to close the achievement gap. The achievement gap is this well known distance and standardized test scores between rich kids and poor kids. It was growing over time. That sounds terrible, but there should be a caveat every time someone mentions this that the achievement gap was growing, but everyone was actually doing better and better at SAT scores and various standardized tests.

Poor kids were doing better over time. Rich kids were doing better over time. Rich kids were doing better and faster over time, so the achievement gap was increasing, but that doesn't mean people were not learning as much. That's just a side. Anyway, achievement gap, so the idea here was to hold teachers accountable. The underlying idea was there are really bad teachers out there, and they're making the achievement gap increase. Let's find the bad teachers, and get rid of them. I don't think anyone would disagree that there are bad teachers out there, and the question of whether they're causing the achievement gap however is much more tenuous, because after all, this achievement gap has always existed in every single country and in every time that we've measured it.

It seems to be correlated to inequality itself. The achievement gap is less there for some unsurprisingly countries, where there is not that much of a gap between rich kids and poor kids. Going back to the goal at hand, the goal at hand was to find the bad teachers, and get rid of them. The first attempt of finding bad teachers was really dumb. The first attempt was this. It was like, "Define a teacher to be bad if a lot of their students didn't do well on the test." That was defined as there were some proficiency score, and if a bunch of their students didn't pass that certain score of proficiency, then that teacher was deemed a bad teacher. Now-

Jacob Morgan: Which at first glance by the way, for most people, they might say, "Oh, that makes sense." Like, "Why wouldn't that make sense, right?"

Cathy O'Neil: Right. It seems like, "Well, they're not teaching their kids the stuff they need to know to pass the test." Well, somewhat true, but remember, we're already acknowledging there is an achievement gap between rich kids and poor kids, so if you put a target on the back of every teacher who has a bunch students that are behind grade level, then what you're essentially doing is singling out teachers of poor kids, because they're much more likely to have kids that aren't up to grade level. In other words, this was a system designed to pick on teachers of poor kids, and it wasn't the point. The point was to improve the achievement gap, to get those poor kids ... get their teachers to be more experienced, to have less churn.

By the way, another thing that's bad about removing teachers is you have new teachers, and new teachers are not as good at teaching as experienced teachers, so you actually want to avoid churn. You want to improve their scores. Anyway, long story short, that was obviously not a very good system for holding teachers accountable. It was replaced by this new system called the Value Added Teacher Model. This is the scoring system that I told you about earlier with my friend who was the principal told me about. The idea was pretty smart. This was the idea. Instead of holding teachers accountable for the final score of their students, we're going to hold the teachers accountable for the difference between the final score of their students and what the students were expected to get.

That's a long sentence, so let's just role play here. I am a teacher, and I have 40 students, or no, I have 25 students is a reasonable number of students. I have 25 students. Each of them comes into my classroom with an expected score at the end of the year. If they're poor kids, you would expect those expected scores to be a little lower, because we already know about the achievement gap. If you're in an affluent suburb, you would expect those expected scores to be higher, and that's what you would have. In other words, I'm not aiming for everyone doing 100%. I'm aiming for these kids to do a little bit better than expected. That sounds like a reasonable thing. The problem is that statistically speaking, it's just not enough data.

There's always 25 kids in my class. What's even more important is that both the expected score and the actual score have a lot of uncertainty around them. The expected score, that model really uses the score that the kid got at the end of the last year. If I'm a fourth grade teacher, it depends largely on what this kid got at the end of third grade. It also depends on how many kids are in this class, how many kids have a free school lunch, which is a proxy for poverty. It depends on which school system we're in, et cetera. It depends on a bunch of things, but it's not a particularly accurate guess. Let me put it this way. If you were asked to estimate one child's score in a year, how likely would it be that you'd actually get the right answer?

I mean, it would be hard. There is uncertainty attached to that number. There's also uncertainty attached to the actual score that a kids got. That doesn't sound right, because you're like, "Well, they only get one score." Imagine if they had taken that test after a night of not sleeping, or after missing breakfast, or on a hot day where they didn't have air conditioner versus a hot day where they did have air conditioner. There's all sorts of things that could change an actual score as well. Long story short, a teacher is being assessed based on 25 kids. The difference is between two numbers, the expected score and the actual score. Those numbers both have uncertainty. Their difference has lots of uncertainty, and there's only 25 of them.

The teacher in question is being ... It ends up being not much better than a random number. The reason I say that, I have evidence for this, the first piece of evidence was I found teachers. I found a teacher who got a six out of 100, and then he got a 96 out of 100. This is the next year. He hadn't changed the way he thought it all. That's not a good sign for a system since you'd think it would vary only if he actually changed his teaching style. Then I actually found this blog post written by a New York City high school teacher who had actually gotten data. It was hard to get by the way. I wasn't able to get the formula. This guy was lucky to be able to get this data.

He found that there were more than 600 teachers who had gotten two scores for the same year for the same subject for different grades. If a teacher was teaching seventh grade and eighth grade for example, and he plotted them on a scatter plot, which is to say if the teachers' scores were consistent, if they got similar scores for seventh grade math and eight grade math, you would expect to see these dots along the diagonal, but instead you saw it everywhere. It was scattered. It was almost uniformly distributed, which is to say that these scores were almost random numbers, and yet, they were being used for this extremely high stakes decisions like getting teachers fired.

Now, you asked me to talk about Sarah. I'm going off on these technical tangents, but I'm going to tell you what happened to Sarah herself, because Sarah was a particularly weird case. She got fired because of her bad overall teacher score. Half of that was due to this value added model score. The other half was due to stuff that doesn't change very much at all. Most of the information was coming from this value added model score. She got a terrible value added model score, but she actually, and this is unusual, actually had a suspicion as to why her score was so bad. This is what her suspicion was. She was a fourth grade teacher. A bunch of the kids coming in from third grade had come on in with very good test scores at the end of third grade, but they couldn't read or write very well.

They were noticeable bad at reading and writing. That's a bad sign. I should say that Sarah was coming in the Washington D.C. school system, where Michelle Rhee was a school chancellor. Michelle Rhee had instituted not only people getting fired for bad scores, but also people getting bonuses for good scores. She started thinking, "Wow, maybe some teacher cheated on these kids' scores." Then she looked into it, and she found out that there an unusual number of erasures from the school that those kids had come from, statistically significant extra number of erasures. There was pretty good evidence that at least she thought that her overall value added model score should be reexamined in the light of the fact that many of her students had come in with what she thought were inflated scores.

Now, think about what that means. If they had inflated expected scores, and she was supposed to be meeting those expectations, but she couldn't possibly meet those expectations if these kids came in not being able to read or write. Cheating the previous year that she thought was happening would have made it really very difficult for her to get a good value added model score. That is the story for Sarah. Now, the last part of the story is she tried to appeal, but she got fired. They wouldn't let her appeal. Their basic thing was, "This is a mathematical algorithm, and it's fair." Again, what really rouse me up is this concept that people think that because it has the word math attached to it, it is inherently fair and objective, and that you have no right to object to it.

Jacob Morgan: It's funny. You actually had a great quote, "Algorithms are opinions embedded in code," which I thought was awesome.

Cathy O'Neil: I say that because when I build algorithms, I make subjective choices all the time. I am projecting my agenda onto my code every single time.

Jacob Morgan: What's an example of how these algorithms and pieces of code are actually subjective, and not as objective as we all think they are?

Cathy O'Neil: Well, give me an algorithm, and I'll tell you. I'll answer that. I mean, the example I like to give is when I cook dinner for my kids, that's the example in my book. I define a dinner to be successful if my kids ate vegetables. My kids would not agree with that definition of success, but that is one of the most important things you do with an algorithm is you define success. Then you optimize to success. If you have a different definition of success, it's a totally different outcome. Another example might be when you are trying to build a machine learning algorithm to higher candidates that are applying to your company, well, you're going to train it on his historical data. You're going to have to say, "Well, who among this current applicant poll looks like someone who is successful in my company in the past?

To define what it looks like to be successful at your company, you could define it to be somebody who got a lot of raises and got a lot of promotions, or you could define it as somebody who had a high productivity level. Those choices really will matter in terms of what that algorithm ends up doing ans who the algorithm thinks will be successful at your company, and who won't.

Jacob Morgan: Maybe we can talk about that for a minute, because I know that's a huge area, what you're hearing about that even specifically in HR, algorithms that help you find the best candidates, algorithm that help you figure out when an employee is likely to leave the company, when they're getting burned, algorithms that help you create experiences for people, and all these companies are trying to figure out, "How do we collect engagement data, performance data, data about the school they came from, and get this picture inside of our organization?" I talked to a lot of HR people, and they're investing so much time and money into people analytics.

I mean, I know that you're not saying it's all a waste of time, or may you are, but how ...? I mean, what do we do? Should we just not bother doing this? It seems like there is no way to avoid some kind of bias or subjectivity.

Cathy O'Neil: That's exactly right. I'm not saying it's a waste of time. I'm saying that it is not objective. It is subjective. If you're going to do it, you have to defend it. You have to say, "This is why I've chosen this." You have to defend it, and you have to make sure that it is having the results that you wanted to have. Going back to the example of productivity as a measurement of success versus getting raises and promotions, we happened to know that most cultures have implicit bias. Most cultures make it easier for tall, white men to get promoted and get raises than for Latina women just as an extreme example. If we define success as somebody who gets lots of promotions and raises, then we're probably tilting the playing field towards the tall white man.

Maybe not. Maybe your internal culture in your business is not like that, but it probably is, and so my argument is that you should make those decisions. You should defend them, and you say, "This is why I'm choosing this." Then you should monitor the results of those decisions, because there is no such thing as objectivity here. There is no real definition. This is the opposite of math. There's no mathematical definition of success within a company. It is up to the person writing the algorithm, or rather, it is up to the person who's going to use the algorithm. You should first think very hard about what that looks like, and whether that's fair. Then second of all, once it's in production, you should make sure it's doing the right thing, that it's actually promoting the people that you think should be successful rather than have historically been successful.

That's the other thing about it is that all of this is basically training on historical data. If we blindly train our culture on historical data from our past culture, then we will never actually evolve. We would just propagate everything we used to be like. Now, if we thought we were perfect, and we didn't need to evolve, that would be fine, but personally, I do think we need to evolve sometimes in a quite meaningful ways. Trusting big data algorithms rather than taking a really hard look at ourselves is not in the direction of evolution.

Jacob Morgan: What would you do? I mean, let's say you're working at a company, and somebody tasked you to help figure out, "How do we create some algorithms to make sure that we bring in the best candidates, or that we have the best programs for our employees?" Is there a best approach? How would we go about doing this to make sure that it is indeed delivering what we want it deliver as opposed to just screwing up the company?

Cathy O'Neil: Well, listen. I think, the most important thing is that algorithms and big data and AI, none of them absolve us for the responsibility of having the difficult conversations. They're not magic bullets. They're not silver bullets I should say. We still don't know what makes a teacher a good teacher even though we had a stupid random number generator that pretended to tell us this. We will never really know what makes a good teacher unless we sit down and decide, and it's no going to be the right answer. It's going to be the thing we compromise on. Similarly, we're never going to all agree on what it means to have an ideal employee. For that matter, maybe we think an ideal employee is good at three different things.

Well, how do you weight those three things against each other? What if they're good at two things, but not so good at the third thing? You have to quantify things to build the algorithm. That's a difficult conversation. It should be a difficult conversation. What I'm trying to say is anytime that you see somebody trying to bypass this difficult conversation using an algorithm, then what they're really doing is they're just shooting in the dark. Who knows what they're actually going to be optimizing too? Having said all that, if you're willing to have that difficult conversation, you'll be willing to say, "Here is what I mean by a good employee. Here is what I want to optimize too," then the algorithms will be very, very good for that, because they are very consistent, and they follow rules.

As long as you know what you actually want, and you write it down, and you monitor the algorithm to make sure you're getting what you asked for, then the algorithm can be a very, very useful jewel for you.

Jacob Morgan: Because with all of these algorithms, we always hear about vendors that offer various algorithm solutions and big data and AI solutions that will help you find the best candidate or ... They always highlight the amazing stories and successes. We improve tenure by 20%. Productivity went up by this amount. Innovation increased. We save money on recruiting costs. When we see those kinds of success metrics and numbers, are we just only getting a part of the picture? Would you look at that from a math background and say, "Okay, well, clearly, this algorithm is working," or would you look at it and say, "Well, that's not the full picture?"

Cathy O'Neil: I would definitely say that's not the full picture. By the way, I should say two things, first that I've found somebody, and I profiled him in my book, somebody who got rejected from one of these algorithms, a personality test, an HR personality test. His father is suing the company that didn't hire him with a class action lawsuit on behalf of everyone who ever took that test on the grounds that it actually valuated the Americans with Disability Act, because this young man, Kyle Beam, got red lighted. He failed this personality test, but he recognized some of the questions of that personality test as being similar to the ones that he was asked at the hospital where he went to for being treated for bipolar disorder.

That's never going to be in a marketing material by a company.

Jacob Morgan: What do you mean he was ...? Meaning that the company knew that he had this disorder, or that ...? What was it like?

Cathy O'Neil: I don't think the company that was using this personality test knew it. It was a third party algorithm. Chronus was the one that developed the algorithm. That's a company in Cambridge Massachusetts. The company that Kyle was applying to was a grocery store called Kroger's Grocery in Atlanta, Georgia. I don't think Kroger's had it investigated this algorithm very much at all. They just had bought into the marketing of Chronus, which is, "Oh, we're going to find you the best job candidates." Little did Kroger know, I assume, that they were actually filtering out people based on this mental health assessment. What I'm trying to say is it's much more complicated than, "Yes, we got a bunch of good employees, because we used this HR tool."

One thing that every company should be aware of is that they will be responsible for illegal hiring practices if the algorithmic tool they're using is in fact illegal. Again, that's one thing I want to say. The second thing I want to say is I actually have started a company, an algorithmic auditing company to try to help companies investigate these tools before they start paying for them, because I don't think companies are sufficiently weary of their own legal responsibilities when they start using algorithms from third party vendors that may or may not actually be legal.

Jacob Morgan: Do you have any advice?

Cathy O'Neil: Does that make sense?

Jacob Morgan: No, no, it does. I'm just thinking of pretty much every company that I've ever talked to is using some kind of algorithm on their people to figure out something.

Cathy O'Neil: Right. It really does matter what that something is. I'm talking about hiring. Hiring is highly regulated so as firing, so as raises. In other words, if they're using an algorithm to do something in a regulated space, they have to make sure their algorithm is legal. People are just completely unaware of this. They just assume that algorithms are inherently legal. It's just not true. It's easy to develop algorithms that are illegal. You could just get rid of all the women in a hiring a algorithm. Now, I'm not saying that's what's happening. I'm not even saying that it is intentional, but I know that algorithms can be unintentionally quite biased and quite discriminatory. We need to wake up to that fact, and start examining and interrogating every single algorithm that we're using in a regulated space.

Jacob Morgan: Have you ever come across any situations, and I'm sure you probably have all sorts of fun stories, where you're told or where you're presented with something, where they said, "You know, this is a great algorithm. It's going to be XYZ," and you looked at the algorithm, and you're like, "What the hell is this?" This is a complete joke. This is how you're evaluating people, or this is how you're doing XYZ. Have you ever come across something like that?

Cathy O'Neil: That's exactly how I felt when I looked into the teacher thing. I'm like, "Are you kidding me? This is a random number. This is a random number, and you're firing people based on it."

Jacob Morgan: That's crazy.

Cathy O'Neil: Yes, it's absolutely crazy. It should not be true. I'll add one more thing about that teacher example. In Houston, a couple months ago, six teachers who had been fired based on their value added model score sued, and a judge agreed with them that their due process rights, their 14th Amendments due process rights had been violated, which is basically that you should be able to understand the system that is assessing you. There is starting to be legal pushback on this kind of opaque secret but powerful algorithm, the usage of this stuff, and these job decisions.

Jacob Morgan: Yes, because, I mean, now, they even have these games that you can play. You download an app. You play a game. Then based on how you play the game, the information goes to this company, where it assesses your soft skills. Are you risk averse? Are you risk free? But you have no idea what it's looking at or why or how. It seemed-

Cathy O'Neil: Look, I'm not a lawyer, but I doubt that's legal. I'll tell you one thing that's already happened with the personality test, the Chronus personality test that I just discussed in the class action lawsuit. It hasn't finished through the courts, but one big barrier has already been passed, which is the question of whether the questions on this test were actually business necessities. Do you need to ask these questions in order to make sure the person, the job candidate could fulfill their duties? The answer was absolutely not. Most of the questions asked were not relevant to the job description. That's a problem. I think, yes, the video game stuff you're talking about, like, is there any legal proof that that stuff is relevant to the job in question? Unless your job is playing video games, I'm guessing the answer is no.

Again, I'm not a lawyer, but it baffles me that people think that any old algorithm that they come up with can be applied in a business setting in a hiring setting. That is a highly regulated space.

Jacob Morgan: What's the difference then between let's say taking a quiz where an algorithm assesses you versus having an individual ask you those exact same questions? I asked that. I know it might seem like a simple question, but if the questions are identical, I mean, somebody is making a subjective assessment, whether it's the human that's asking you those questions, or whether it's the algorithm that was created by a human. Let's say-

Cathy O'Neil: The answer is there is no real difference, and also, there is lots of case law. Again, I'm not an expert on this stuff, but I know there was an example of, I mean, I want to say, firetrucks. I mean, not firetrucks, not the police but the firemen where they were starting to give everyone an IQ test or something along those lines. It was in order to avoid hiring African-Americans, and it was deemed illegal. The reason was because you didn't actually have to pass this test to do your job. There is a general effect about hiring, which is that you're not supposed to be asking irrelevant questions just so that you can filter out people you don't like.

Now, having said all that, most companies, when they ask irrelevant questions, they get away with it because they haven't systematized it. That's exactly what you are doing when you develop these algorithms. You're systematizing these questions so that they're an artifact of all the things you've been doing with hiring. In that sense, I will say that it's actually good, that we're using algorithms for this stuff, because at least now, we know what we're doing wrong. If you think about it in 10 years, maybe we will have actually achieved something really, really good, a progress with hiring algorithms namely that we will have developed systems of questions that are relevant and important, and nondiscriminatory and legal.

That would be an improvement over our current system.

Jacob Morgan: I'm totally shocked [inaudible 00:48:44] that you're saying that, because so much of our lives are run in algorithms, pretty much everything from Amazon to Netflix. I have a Facebook page, where I buy Facebook ads, and it does a lot of those assumptions that you make. You pick a certain target, and it says, "Oh, well, these people are likely ... This is their household income." You're making me wonder in a world that's being run by algorithms pretty much everywhere, is everything just like a complete façade? Are we like in the matrix, and everything is just going to collapse one day?

Cathy O'Neil: Well, I do think based on my own experience of picking the winners versus losers that the algorithmic culture that we now live in is a direct threat to the American dream. That's not an understatement. I mean, the American dream is that somebody starting with nothing can be a success, but right now, what these algorithm are doing to us is that they are deepening the standard paths that we are expected to down. Oh, you're from this zip code, and you're this gender, or this demographic. Here is what we expect from you, and here is what we're going to make happen. In other words, we're not just predicting the future. We're making the future. When we make the future, and we keep people who are unlucky, keep them unlucky, make them even more unlucky.

We keep people who are lucky, lucky, and make them even luckier. That is not social mobility. That's not the American dream. That's where we're going unless we take a U turn, and start saying, "Hey, wait a second, this isn't fair. This isn't fair to say just because you happened to be living in the wrong neighborhood that you should never get a job." That doesn't make sense. We have to push back very deliberately. That's what I mean by the evolved society. We have to ask for more than what's happened in the past.

Jacob Morgan: Is there a difference between good algorithms and bad algorithms, because we have algorithms that control traffic lights? We have algorithms that control energy use. Then we also have algorithms that control the types of ads that we're presented with. I'm trying to figure out is there a good algorithm versus a bad algorithm, or is every algorithm that's created in some way bias, and we just need to accept that it has some subjectivity and bias to it?

Cathy O'Neil: Well, every algorithm is subjective. I care of course much more about algorithms that give or take away options for people's lives than I care about traffic lights. I make a very important distinction in my book that I'm not talking about all algorithms. I'm talking about important algorithms that affect a lot of people in important ways, so they have to have rights to a certain level of impact and influence before I even think about them. When they do, then yes, we cannot strive for objectivity since there is no objectivity in most of these domains. What we have to strive for is something that we are comfortable with ethically, morally, and legally. That would take some doing, because right now, that's not how we approach machine learning. How we approach machine learning is throw a bunch of historical data at the wall, and see what sticks. Then assume that it's perfect.

Jacob Morgan: What do we do for people that are listening to this? Well, I guess, there are two groups of people that might apply to this. Either people who have had to deal with an algorithm then made some sort of a life decision for them, maybe it was a loan that they didn't get approved for, a job that they didn't get hired for. First question is what advice do you have for those people who are constantly perhaps being eliminated in their opportunities by algorithms? The second is what advice would you give for organizations that are themselves designing these algorithms, and turning them inwards on their people to figure out who they should promote, or even algorithms that are used to assess customers?

Cathy O'Neil: For the first group of people, I would say that if you're lucky enough to know that you're being scored by an algorithm, which most people actually aren't, these things happen invisibly often, but if you are lucky enough to know that you're being scored by an algorithm, then absolutely demand explanation. Please do make use of the fact that the 14th Amendment says that you do have due process. I really think that this is a political fight. This is not a math test. You don't have to be a math expert to say, "I deserve to know why I'm being fired, or I deserve to know why I missed that promotion." You deserve it because you're a human being. That's the first group of people. The second group of people I'd say is like, "Dude, audit your algorithm. Hire me to help you."

I know that sounds very commercial, but I am currently developing tools and researching tools for how to keep track of these black box algorithms to make sure that they're doing what we want them to do. I would love to teach people how to think about this, and how to design algorithms from the get go to work better than random, which is what we have now.

Jacob Morgan: Yes, random is not good. Well, I know we're just at 2:30. I don't know if you have a few more minutes or if you need to jump off right away, but I want to be respective of your time. If you have maybe a couple more minutes, great. If not, totally understand it.

Cathy O'Neil: I really do have to go, I'm so sorry.

Jacob Morgan: No worries at all. You've been very gracious with your time. Where can people go to learn more about you and your book, and even this practice that you've built?

Cathy O'Neil: Well, my company's at oneilrisk.com, O-N-E-I-L risk.com. I write a blog about at mathbabe.org, which is sometimes about algorithms, sometimes just about my crazy life, and my blue hair, and my recent bariatric surgery. I also am a Bloomberg contributor, do if you ever go to bloomberg.com, you can look me up there. I've written a lot in the last few months about algorithms and big data. Some of them are optimistic. Some of them are less.

Jacob Morgan: I definitely want to encourage everyone to grab a copy of the book. We've just touched on the very, very surface algorithms and data. The book explores this in much, much more detail. Cathy, thank you very much for taking time out of your day to come speak with me.

Cathy O'Neil: It's my pleasure, and thanks so much for having me.

Jacob Morgan: My pleasure. Thanks everyone for tuning in. My guest, again, has been Cathy O'Neil, a mathematician, data scientist, and author of the new book Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. I'll see all of you guys next week.