EPISODE 1526

[EPISODE]

[0:00:00] J: Do you want to start this off just with a quick overview of your company privacy dynamics and brief history? Then what you do and yeah, we'll take it from there.

[0:00:14] JC: Sure. Yeah. We're Privacy Dynamics. We're a seed stage startup. We've been around for three years. Our vision all along has been to unlock data that's not being used today. There's a lot of data. It's not used for fear of re-identifying individuals. There's a lot of data that's not used, because it's very subjective as to how risky a data set is. If you have a data set with healthcare information, or consumer data, or financial data, a compliance officer, which they are typically lawyers, they'll say, "I can't tell you how risky – what a risky data set is, but I'll know it when I see it."

That's the underpinning of the company was to come up with a way to quantify data risk in a repeatable, mathematical way. Along the way, we realized that if we knew what was making a dataset risky, we then also knew how to anonymize it and make it safe. Really, that's what we're trying to do is we're trying to make it possible for data to be shared within large organizations, between one org and another, where data can't be shared, or between different companies.

Speaking of anecdotes, I don't know if this is still the case, but up until not too long ago, it could still be true. I'm here in Seattle. The Seattle Fire Department wasn't sharing data with the Seattle Police Department. It sounds crazy, when you tell someone that doesn't know much about data privacy, do we tell them that? It blows their mind. When you think about things, like a 911 phone call database, that actually contains a lot of sensitive information.

Graham, our CEO, he was at Microsoft for seven or eight years, and he was working in Azure at the time, and this is a while ago, but trying to get government agencies, municipalities, private companies to move to the cloud. The hang up was really moving their data to the cloud. He saw a big opportunity there for changing things and disrupting things in data privacy.

[0:03:18] J: The data that, for example, you want to move to the cloud, of the data that could be shared between a police department and the corresponding fire department, in both instances, that would have to be the real live data of it, right? In these specific examples, anonymized data would not be what they want. If the fire department shares an anonymized version of its data with the police department, or vice versa, would that be useful?

[0:04:04] JC: Yeah. I think, if what you're getting at is data needs to be shared in real-time, versus data that can have a delay – yeah, I mean, there are lots of different applications in data privacy. There a lot of different companies in the space and there are a lot of different problems that need to be solved. Where we really started was not really data, not real-time data, but if you need to do some analytics on some data and you're okay with a five-minute lag, or a 24-hour lag, or even a 30-day lag, that was really where we started.

If you're trying to do some analysis of say, 911 phone calls to figure out what times a day are the calls being made? What types of cases are being reported, like those sorts of things, well, you need to share that data with the data science team. They don't have to have it in real-time. There are other companies – I'm sure there are other technologies that can solve the real-time challenge. We saw that as a just a small piece. A bigger piece is just analytical data.

[0:05:46] J: Yeah. I think there's a bit of lag. I came across that use case of data analytics as well and another in my career. Then another one, I think, your website talks about, which is just data in other environments, like staging environments, developing environments, you still want data that to the system feels like the actual data, so you can trust the load test to run, or so you can show demos that actually make sense. Obviously, you don't want to have actual sensitive client data in the lower environments. On your website, you also speak of deidentifying. Is this the same as anonymizing? Or is there a nuance?

[0:06:35] JC: Deidentification is really more of a formal term to describe removing PII. Anonymization is going a step further. Or the way we look at anonymization is greatly reducing to a very, very, very small chance of an individual being reidentified. There is legislation here in the US and all around the world. The big legislation here in the US is HIPAA, CCPA, CPRA. They spell out very prescriptively what the identification is. It's a little bit different, I think.

[0:07:35] J: Interesting. Interesting. What would you say the challenges are in either deidentifying, or anonymizing data? Obviously, don't spill your whole company secret sauce. I suppose, a good way of framing it would be if naively now I said to myself, I have a small company and I want to write my own quick, let's say, Python script to deidentify data, what are the quick pitfalls I'm going to fall in very quickly?

[0:08:13] JC: Sure, yeah. I would caveat that by saying, it depends on the industry you're in and the type of data you're working with. Yeah, but at a high-level, people are – when they think about deidentifying data, they're just talking about removing direct identifiers. That obviously, those need to be taken care of. Either removed, or masked, or replaced with synthetic data. There also are indirect identifiers. Those are things like age, gender, marital status, nationality. Those things by themselves are perfectly fine, but they can be combined together to uniquely identify someone. That's really what we do is we make sure that we go above and beyond just removing direct identifiers. We also take care of the indirect identifiers and make sure that individuals can't be reidentified.

[0:09:37] J: Yeah, absolutely. That makes a lot of sense. Not always intuitively, if you think about the most common cases, but I suppose, it's especially in edge cases, where incidental reidentification can occur. I thought about this a few times when I was at uni, for example. If you ask someone their nationality and the first name like, “Oh, it's John from the UK, at University of York,” that doesn't really narrow it down. If I told people I was Jeff from Luxembourg, those three pieces of information that I studied at that uni, my country, which obviously, very small, and then my first name, because there was another guy from Luxembourg at that uni, would actually identify me precisely. Yeah, it doesn't always apply to the first big categories you think of, but then it would in smaller cases.

How do you ensure data cannot be reidentified? Do you run some simulation on the deidentified data to see if the machine’s best guesses can actually yield some results that would jeopardize privacy? Or how would you go about that? Is it more, you have a model that's mathematically already proven to deanonymize and you trust in that? How much I suppose, how much upfront theoretical work and how much dynamic computational work goes into your approach?

[0:11:24] JC: Well, it starts with proven methods. There's a lot of academic literature out there. I mean, there's just mountains and mountains and mountains of it. Our approach is based on K-anonymity. It's not exactly we have our own flavor of it. I think, it starts by using proven methods. We also work with the foremost expert in deidentification, who wrote the guidance for HIPAA. We started with healthcare data, which is probably the most rigorous –

[0:12:06] J: Heavily protected and, yeah.

[0:12:09] JC: Yea. It's the hardest data. It's most important data. Yeah. Then we run a statistical attack model, making very pessimistic assumptions to identify the weak records. We assume that the attacker has a lot of really good data. Then we look at the weak records in the data, and we blend them in with the other records by perturbing the data, so that you can't identify a single individual.

[0:12:51] J: Interesting. Would it be an accurate analogy to say, going back to my previous example, self-centered, though. Maybe if there is an entry in a database that says, “Oh, this is someone called John from Britain,” we're happy to have that as their category. If there's someone called Michael from Andorra, we'll put them in the same, as in we won't have Andorra as a category, but we'll aggregated with loads of other small countries, so that on average, there might be more than one Michael from that group of countries. Is that the analogy here?

[0:13:36] JC: Yeah, I think so. My knowledge of European geography isn't fantastic. I know where Andorra is. Actually, Andorra has been on my list of places to go for a long time. Another thing that we might do is we might say, they're not in Andorra, they're in southern Europe, or they're somewhere between France and Spain. If there was a larger region that you could move them into that would be more obscure, then we would do that.

[0:14:17] J: Yeah, yeah. I think I see. You talked about K-anonymity just now. Can you briefly explain what K-anonymity is, or used to know?

[0:14:32] JC: It's essentially bucketing records into a certain size. K is the size of the bucket. You could have K equals two. That means you've got at least two records that look the same, or you could have K equals 10, or K equals 25, K equals 50. If you've got K equals 25, then those

25 records are – they're all going to look the same. That's how a lot of the data brokers work is they're sharing data with bucket size of 25, or 50, or something like that.

[0:15:22] J: Yeah, that makes sense. You're basically putting down the resolution of data from one, where you can see each entry individually to a lower resolution, where you can see 25 at same time. I suppose, that would make it impossible to, if we go to the example of analytics, would make it impossible to run analytics on outliers, of which there is only ever one. But that's unavoidable, I suppose. There is no other methods of anonymizing what you can still take into account outliers in your analysis.

[0:16:11] JC: Yeah, we're trying to make sure that no one record can be identified. Sometimes you might have some extreme outlier. I think, there are different ways to address that. Making sure you don't have the direct identifiers, and also, knowing what some of those outlier values are, but not with them all combined together. That's less of a concern for us. Really, what we're trying to do is get more data into the hands of people to do their work that otherwise wouldn't be able to get access to that data. I feel like, if we can get someone 90% of the way the data that they need, that's better than zero, right?

[0:17:16] J: Oh, yeah. Absolutely. Can you name a few of your most prominent, or most representative use cases, and what they're doing with it, or the approach that they're taking? I think I might give listeners a more concrete idea.

[0:17:37] JC: Sure. we started with data for analytics. Really, our algorithm is really fast, so we can run – we run in batch mode, but it's very fast. You don't have to wait very long to process data. As a result, we are plugged into an ETL, or ETL pipeline. Data's landing in a data warehouse. It’s going through some transformations. One of the final steps will be to anonymize the data. Then that way, you have some schema, or somewhere else that you can give a larger number of people access to to do their work.

What we've seen more recently, is this uptick, extra traction with anonymizing data for just really, just more traditional software engineers that need data for development and testing environments, especially as more and more software development is moving to the cloud with code spaces and – I mean, there's a whole host of these developer preview environment,

companies that are more or less based off of Kubernetes. But they still need data, right? You can have all your application code running, but you need realistic data. Ideally, you're pulling it from production.

What we're actually seeing people do now, we're seeing them take a – anonymize a copy of production database on a daily basis, or maybe a couple times a day and give that to the engineers.

[0:19:34] J: Yeah. Both of those use cases are batch anonymizing. Do you do any real-time yet? Are you looking into it? Or are you deciding actually, no, we're going to specialize in what we're really good at and do it patch. What's your approach in regards to that?

[0:19:58] JC: No, we'll definitely do real-time data. We'll definitely do it. We're not quite there yet. Yeah, I mean, it's definitely something we're working on and something we'll have.

[0:20:11] J: Does that pose any more challenges with regards to the K-anonymity approach? Because you might not know what the buckets should look like at the beginning, or bucket split might change over time. Or is it actually pretty straightforward to go from batch to real-time?

[0:20:37] JC: My engineering team would probably kill me if I said that it was straightforward to do that. It's not. We have some pretty good ideas on how to solve that problem. I mean, the reality is even just batch processing, I mean, we're a relatively small company. Even just batch processing will keep us busy for quite a while. We do see longer-term that there's a definite need for streaming data.

[0:21:17] J: Yeah. In a similar vein, to not having access to all of the data at the same time, I was wondering, is there a security concern? I like to think I have a bit of a security and privacy minded brain anyway, so maybe this is just too far and too technical. Is there a concern that if you have a batch of anonymized data, a potential attacker with access to previous data, also anonymized, but as in historical versions of that same data might be able to reconstruct some of the original data that would not have been possible to construct with only one snapshot and time off the data and therefore, that your system wouldn't pick up on being able to reconstruct? Am I making any sense?

[0:22:15] JC: You make perfect sense. Yeah. I mean, there are there are definitely examples of data that was released, thought to be anonymized, and was later proven not to be anonymized. I think the Netflix example is a good one.

[0:22:43] J: Can you recap for us quickly, what the Netflix example was?

[0:22:48] JC: Yeah. It's been a while. My memory is fuzzy on this afternoon. Essentially, they – I don't know if you recall, but they had this programming algorithmic challenge to come up with a better recommendation system. They released all of their movie ratings as a database. I don't remember, if it was a CSV, or I don't remember exactly. They took out the direct identifiers. Then later, some people were able to combine that data set with other datasets and re-identify individuals. Again, that speaks to the importance of perturbing the quasi-identifiers, or indirect identifiers. That's one example. We're doing these simulations in real-time, over and over and over and over and over making assumptions, very pessimistic assumptions about all the other –

[0:24:07] J: Yeah, I need to say, you assume that an attacker would have a lot of really good data already.

[0:24:13] JC: Yeah. I think that's the key is – Some cases, you might have to have a K size of something larger than two. I think, by being very pessimistic about what an attacker might have, gives us a lot of confidence in reducing the risk of re-identification.

[0:24:52] J: Yea. Just to reiterate, a case size of two, would that mean that an attacker would have a 50%-50% chance of guessing which individual specific record is talking about? Then I suppose, legislation like HIPAA would say, K has to be at least 50, or something like that. Is that correct?

[0:25:22] JC: Yeah. HIPAA doesn't spell that out. HIPAA is – that is just not something they define. They spell out all of the direct identifiers that you have to take care of. I don't remember the number. It's 15 to 20 direct identifiers, or maybe even more. It bids in that neighborhood. They also define two different ways of sharing data. One is called – is referred to casually as expert determination. That's when you have someone well versed in statistics and privacy risk,

that does an assessment of the data. They've determined that the risk of reidentification is very, very small. That is typically done on a case-by-case basis, but that's really – that was how we got started with healthcare data was doing that in an automated fashion.

[0:26:39] J: Yeah, because that sounds quite labor intensive, if you have to have someone highly qualified, assess each data set individually. Do you have a use case for supporting data that HIPAA, do you remember?

[0:26:58] JC: Safe Harbor. Yeah, it's called Safe Harbor. It greatly reduces the utility of data. I don't remember exactly, but it basically it spells out and says, if you have a US zip code, you have to remove the trailing three characters, or something like that. If you're doing any analysis on this, where you're trying to deal with locations, it's going to give you some very rough data to work with, and it will be hard to derive any insight from it.

[0:27:39] J: Yeah, I see. Do you think it would be possible to abide by the letter of HIPAA, removing all the direct identifiers and such and still have a data set that is really identifiable through other means, like secondary identifiers and so on?

[0:28:10] JC: It's a very open-ended question. Yeah, I would have to look again at what Safe Harbor is. What are they're spelling out to obscure. Again, with expert determination, it's pretty specific about the use case. Yeah, well, it's also, I don't remember the language exactly. It's data being used for this use case by these individuals. It cannot be reidentified. Individuals in that dataset cannot be reidentified. That's different than data that's been anonymized and is shared with an audience that it wasn't intended to be shared with. Again, that's why I think being as pessimistic about what the attacker might have access to is really important.

[0:29:21] J: Absolutely. There is not just HIPAA, of course. There are a number of standards, a number of different legislations. Then there are use cases, like the Netflix one, where you share the data with the wider public. Obviously, you don't even have the protection of you extensively trust the recipients of the data, but it's just as soon as you release something into the Internet, anything goes as we know.

There are some really cool encryption technologies out there. One I've come across is convergent encryption, I believe it's called, which allows you to have searchable ciphertext. Your data is encrypted. I know what you're doing isn't specifically encryption, which is as an example, you can have encrypted data that you can still search through. Then you still have, you can find the result, but you still have to decrypt it, to see what it says. Which is just mind-blowing to me, if you think about it. What is the most mind-blowing to you, piece of technology you've come across in this domain? Are you using it? Or does it not have a place yet in Privacy Dynamics?

[0:30:42] JC: Well, I think, the encryption example that you came across, I think it's – I do agree. It's very impressive. Well, first of all, we don't do encryption, but at the end of the day, the encryption is solving a data access problem. Even if you have this encrypted data, if you're restricting people to aggregate queries, and if you're not using something like K-anonymity, or differential privacy, you can still leak sensitive information. I think there's sometimes when the encryption makes sense, but it's also not a silver bullet. You still have to ensure – That's what we're trying to do is eliminate the chance of human error.

We see this a lot with compliance software. There are all these specific rules that are set up, but for data access. Ultimately, those rules are defined by humans, which are certainly capable of making mistakes and exposing some data that shouldn't be exposed, even though you think it's protected. What we're trying to do is just say, okay, hey you run your data, we anonymize it down to a certain K size, and it lands in this data warehouse, or this relational database that your engineers are going to use. It's safe to use. You can give it to your entire engineering team, or you can give it to your entire data science team, or BI team.

[0:33:03] J: What you're seeing is encryption addresses data access, as in someone who isn't authorized to shouldn't have access to the data in the first place. You're not tackling that problem. You're sidestepping it, in order to make the data okay to be accessed by a wider group of people, or even anybody.

[0:33:30] JC: Yeah. I think, you summarize that very well. Yeah. That was great.

[0:33:37] J: Going back to the question about just any cool technology you've come across and not focusing on encryption, is there anything that is part of just research, which I suppose

assume that you do quite a bit of at Privacy Dynamics, anything that you found that someone has proven always working on on any tech, or algorithm, or mathematical model, or whatever it may be that just personally blew your mind?

[0:34:14] JC: This is really outside of our domain, but I'm really excited about all these efficiencies that are coming to software development in general. I think that these generative texts – What's the what's the GitHub product? Copilot. Copilot being combined with developer and preview environments in the cloud. All these really advanced tooling to make software engineers – I mean, really 5, or 10X more efficient. I think there's a lot of momentum in that space.

I think there's a future with Privacy Dynamics there in that space, unlocking data. Just someone that's been doing this and writing software for a really, really long time. I'm pretty excited about all that. I think that's going to have a big impact. A really, really large impact, right? Because there's so many software engineers around the world.

[0:35:45] J: Yeah. I think it's a really interesting thing to look at how the priorities have shifted over time. Back in the day, you would have been more focused on writing code that maybe fits into very little memory. I don't know any of the numbers, but sometimes people bring up how little memory the rocket used that went to the moon. The programming that had to be done in there had to be super-efficient. Nowadays, we don't care about that at all. We care a bit about network throughput, but mostly we just care about creating an MVP, minimum viable product for our company to prove its viability. We care about the time of the engineers, not the time off – or the CPU.

I can absolutely see how, if I wanted to set up a company, I just wanted to have some stock development model and environment. I'll use a cloud that's well-established, even if it's not the cheapest, because I care more about the cost of my development team than about that of the cloud resources. I'll use some other software and collaboration tools. I think, you're absolutely right, there's definitely a place in that stack for Privacy Dynamics, so that I don't waste any of my own precious time coming up with data for my other environments, or for my analytics that either takes me way too long, or is useless. Or thirdly and more dangerously, is useless and I don't even realize it. Yeah, I absolutely get what you mean.

[0:37:38] JC: I think there's going to be a big change in how software engineers work. We've worked the same way for a long time. I mean, I'm guilty of it, too. I still run off of my laptop, right? Because what if I'm on an airplane and I need to do some work? Really how many hours per month do I actually spent on an airplane? Not that many, right? I think, we're going to see a big change here very soon. Another thing too, is with these preview environments that I think are cool is you can spin up in a huge stack that you could never fit on your laptop.

[0:38:29] J: I get what you mean. I still think there's a case to be made for small feedback cycles. I get the airplane analogy as well, because on my current project, I've ordered a laptop, rather than a VDI, because I knew I was going to fly into London and back as I did last week. Truth be told, I was too hungover to do any work on the plane anyways. It was more of an idealized scenario than one I actually needed to draw back on. Yeah, I still think there's a case to be made for quick feedback loops. I want to be able to run, develop stuff on my laptop and just quickly say, try this, try this, try this and then I'm happy. Then I'll bring in a bigger environment.

I think if anything, that's an even stronger case for something like privacy dynamics, because it's one thing to have poorly anonymized data, or data I've anonymized myself, in my own environment. It's another thing to put it on cuffs, on developers' laptops and send them around the globe on planes hungover.

[0:39:43] JC: Right. I mean, let's be honest, there's no engineer that really wants to spend their time writing scripts to convert the production database into something that can be used by the rest of the engineering team. I mean, so far, we've seen a lot of uptick for us there, just because with a few clicks, you can have an anonymized database that's useful for everyone and it's super easy. I mean, just like, there's some things that we just don't do ourselves. We outsource to other products, or services, right?

We'll use another service for source control. We're not going to run our own source control server and have to maintain it. We're not going to run our own data center, right? That would be too expensive and –

[0:40:50] J: That's what I meant before. If I try and do it myself, I'll do it poorly. I'll do it poorly and I won't even realize I've done it poorly. I'll rely on poorly representative data for my analytics or something, and I'll spend way too much time on it. I'm strongly in favor of outsourcing these things to a much more experienced and therefore, competent team also.

Speaking of teams, we've talked about the company history a bit at the beginning. I'd like to touch a bit on just the work-life balance as well. You talk about that on your website as well. First, you also talked about enabling ethical data teams. What do you mean by this? Why is this so important to your company?

[0:41:51] JC: Well, the main point is what's often referred to as data minimization. We want to make it – What that means is that individuals that have access to data only have access to the data that they need to do their jobs, and only for the duration of time that they're doing that job or project. It's coming. There's some new legislation that's taking effect in 2023, the CPRA, the California Privacy Rights Act. It's being more prescriptive about data minimization. One, we think it's just the right thing to do. More importantly, or equally as important, companies are going to have to start thinking that way. Really, what we're trying to do is just make it as easy as possible for a company to do the right thing there with data. Anonymize it and help you only use it – hold on to it for the amount of time that you need it and no longer.

[0:43:11] J: Yeah, absolutely. I think, to summarize this, just to make sure I've understood it. Ethical data teams, or ethical data consumption is only consuming the data that you need. If I sign up to YouTube, let's say, it makes sense for them to know data, such as my date of birth, in order to know which videos they can show me, but they don't need to know data, such as I don't know, my hair color and my shoe size. Enabling teams to see only the data that they actually need to see for their job, and therefore, eliminating any opportunity for unethical data consumption. The first example that comes to mind is, I think, a bit too extreme. Thinking back to Snowden and how NSA people used to look at data, a bit beyond their remit, their friends, girlfriends, and so on. Basically, eliminating that opportunity altogether, and therefore, ensuring ethical consumption. Is that about, correct?

[0:44:29] JC: Yeah. Once again, you did a really excellent job of summarizing and actually doing a better job of explaining it than I did, I believe.

[0:44:42] J: I'm a professional podcast host after all, even though it is my first time. Thank you very much. That's a lovely thing to say. Tell me about just working at Privacy Dynamics. You do emphasize work-life balance on your website. It means different things to different people, I've come to find out through well, talking to different people. What does it mean to you? Why is it important? How do you achieve it?

[0:45:16] JC: What does it mean to me? It means that you have a place outside of work. You have somewhere that you can go with your brain that's away from work. You can't think about work 24 hours a day, seven days a week. I mean, everyone's different. Some people spend more hours per day than others. I think the important thing is that you take a break, and that there's some balance there. Each person needs maybe slightly different balance, but if you don't stop and recharge, then you're going to burn yourself out.

I run a lot. It's like running a marathon. You can't go out, running as fast as you possibly can, you won't make it to the finish line. You'll simply run out of energy. You have to pace yourself. You can't pace yourself without having some sort of balance with –

[0:46:41] J: Yeah, absolutely. Me personally, and obviously, I think a lot of people who work in the tech industry, in the startup industry, a lot of our listeners will be familiar with that. I myself have had a bit of a burnout depressive bout. I had to stop working quite early on in my career. What would be your approach to limiting that? Because like you say, everyone is different? Does it just come down to talking openly about it in the workplace, being a bit more mental health aware and making it okay for people to express how they feel? Or do you have a more global approach? How would you go about that?

[0:47:30] JC: We're still a small team. We don't quite have a global approach. Yeah, I mean, it makes it easier. I mean, I think we emphasize stepping away from work. We emphasize that across the team. There's so many different ways to get it across, but we emphasize it to everyone. We encourage someone. Again, here in Seattle, it's gray and gloomy through the entire winter. If there's a day in the middle of the winter, when the sun comes out and they're blue skies, but it's 1:00, well just stop what you're doing. Your work can wait, and just go enjoy the weather.

Yeah, so it's those things. Everyone's encouraged to take those sorts of breaks. I think, there even are smaller little details that you can do. Don't send emails on a Saturday. I'm a big fan of scheduling emails, or Slack messages. If something comes to my mind and I'm thinking about it and I want to write it down, I'll go ahead and write it down. I'm not going to send someone an email on a Sunday afternoon, asking them to work on something. I think that's part of it, too, is being respectful of – being respectful of other people's times.

Everyone has – they have different interests in what they do. We have a woman on the team. She's a comedian. She does performances. I may be busy doing my non-work things at different times of day, than she is doing her non-work thing. I think, also, being aware of not everyone is necessarily like you, and they break up their time in different blocks than you do. Being respectful of that as well.

[0:50:10] J: It sounds like, it's a lot focused around open communication and trust, which I think is one of the advantages you have as a smaller company. I've seen companies. It's painful, but I've seen them scale that to still be the case when they have hundreds of employees. I hope that continues to occur for you. On a last note, what's your approach to remote working? Do you encourage it? Do you encourage people to come into the office? Is it up to everyone themselves?

[0:50:46] JC: Most of our team is remote. We have a small team here in Seattle, but we have people across North America. Yeah, I think, remote work is great. You do need to have some way of being productive when you are working. You also need to have – within an organization, you need to have some agreed upon time period when everyone can be online together. I mean, I think, remote work is here to stay. The exact definition of it may change a little bit. I think, after 2020, I think the world has changed. We get everyone together a few times a year. When we do get together, there's no agenda, or prescribed things, prescribed work things that we have to do. It's just time to spend together and bond and get to know each other in ways that's difficult to do over a video call.

[0:52:06] J: Yeah, absolutely. I think that's a great idea. I try and encourage my colleagues and other co-workers, that I may have to do that as well, because you just react completely

differently to say, a Slack message, or an email, if you have seen the person face-to-face than if it's just, “Oh, that guy again.” Yeah, I can absolutely empathize with that.

Cool. Well, John, thank you very much for coming on the show. Don't really know how to sign these off. It's fine, we can cut it. Do you have anything else you wanted to add? Anything else you wanted to talk about?

[0:52:52] JC: I think, there's one thing that I would still add, going back to the work-life balance. I just keep it very succinct. Not every company can do this. We've had the advantage of growing at a very consistent, steady clip, but not any really large aggressive spikes. We want to continue to do that. I think, by not absorbing, or scaling up in large, massive growth spurts, then I think that makes it easier to maintain that work-life balance that you spoke of.

[0:53:50] J: Yeah. I think, I know what you mean. If you bring in too many people at once who aren't familiar with the culture yet, it's going to irretrievably dilute that culture. But if you bring them in slowly, one by one, they all – I have a semi fleshed out analogy in my mind with some balls dropped in colored liquid and then they can acquire the color. Then if you do it slowly, then they'll spread it to others. It doesn't really work yet, but I'll keep thinking. Yeah, that certainly does sound like an advantage in keeping a company mentally healthy.

[0:54:33] JC: Well said, Jeff.

[0:54:35] J: Are you hiring, or are you looking – Is there anything else I should mention? Like, apply here or something like that?

[0:54:44] JC: I just tell people to go to our website, privacydynamics.io. If we can make your life easier as a data scientist, a data analyst, or a software engineer, we'd love to.

[0:55:01] J: All right, so that brings us to the end of the show. Thank you very much, John, for coming on. If you're interested in learning more about Privacy Dynamics, go to their website, privacy, or privacydynamics.io. Feel free to get in touch with them.

[0:55:22] JC: Thanks, Jeff. Take care.

[0:55:25] J: Cool. All right –

[END]