Transcript of ‘AI as a Force for Good in Modern Slavery Reporting’

Season 3, Episode 42, Transforming Tomorrow

[Theme music]

Paul: Hello, and welcome to Transforming Tomorrow from the Pentland Centre for Sustainability and Business. I'm Paul Turner.

Jan: And I'm Professor Jan Bebbington.

[Theme music]

Jan: Love it or hate it, find out how AI is being used to analyse companies' modern slavery statements.

[Theme music]

Paul: Jan, have you ever faked a reference?

Jan: What do you mean, faked a reference?

Paul: Have you ever written something about yourself, uh, in the hope of getting a job, pretending to be someone else?

Jan: [indignantly] No! [laughs]

Paul: I'm going to read some words to you.

Jan: Heavens.

Paul: I have three different sentences.

Jan: Uh-huh.

Paul: Jan Bebbington is a prominent academic known for her work in accounting, sustainability, and social and environmental reporting. She's widely recognised for advancing how organisations measure and report their environmental and social impacts.

Did you not write that about yourself? Or is that quite complimentary enough?

Jan: No, it sounds a wee bit flowery.

Paul: Okay. Right, right…

Jan: …yeah, yeah.

Paul: Jan Bebbington is important because she helps redefine what accountability means in business and government.

Jan: Oh. I don't like the important thing, but the other bit…

Paul: …oh, okay, right…

Paul: …is not so bad.

Now, now what about this one? Professor Jan Bebbington is a truly remarkable scholar [Jan laughs] whose career exemplifies the very best of purposeful, impactful academic work.

Jan: Well, I, we, we, I'm not gonna say the word, but that's BS, I would say. [laughs]

Paul: So you're trying to tell me you didn't write any of these?

Jan: [laughing] No, I don't think I did…

Paul: …'cause from what I'm told, AI has to have a source to go on when it creates, um, words and profiles…

Jan: …well, all of those words exist…

Paul: …so when I typed into it, [Jan laughs] 'can you describe Jan Bebbington?' and it came out with these three descriptions, you had nothing to do with it. You haven't been, you know, fiddling the system?

Jan: No, no, I, I can tell you I haven't. And I tell you why I think it's AI is that it's both too…

Paul: …'cause I just told you…

Jan: …no, too general, but also too fluffy.

I mean, that would be, that would be like, toe curling to describe yourself in that way.

Paul: I, I'm going to read another one. And you have to guess who this is about.

'X is, by any measure, a quietly indispensable force at one of the UK's finest business schools.'

Jan: Do you think it's our guest or is it you?

Paul: It's me. [Jan laughs]

I'm glad you didn't think it was you, 'cause I did, I, I left that, the name out of that one, and you might have said, what was that me as well? But no, no. Apparently, I'm a quietly indispensable force.

Jan: I find that very uh, I find some of that description, uh, inaccurate.

Paul: But, I won't read the name of our guest yet, but this is what it says about our guest.

'He is a scholar of genuine distinction, bringing intellectual rigour and real world relevance to everything he touches.'

I particularly like the, 'everything he touches' [laughing] bit at the end. [Jan laughs] I like to think that when he's at home making dinner, you know, he's bringing real world relevance to… [tries to suppress laughter and fails]

Jan: I'm sure, I'm sure that is the case.

Paul: …to, to his meals.

Um, we've got an episode about AI, Jan. Had you guessed that?

Jan: I had guessed that. But um, also, I, I, hope by the end of it, 'cause I have a, a bit of a sense of what we're gonna talk about, that you'll realise there's different uses of AI and what you've just described as maybe, uh, is entertaining though it is, is not, perhaps the best use of AI.

Paul: No, maybe not the best. And it irked me somewhat because in order to get these answers from two different AI engines, I had to create an account for one of them. One of them I have via work because it's done by a company that does just about everything…

Jan: …oh, yeah…

Paul: … on your computers, but the other one I had to create an account just to get it to tell me what…

Jan: …I love the word 'irked'…

Paul: …yes… [laughs]

Jan: …I hope this isn't indicative of your mood for the entire podcast…

Paul: Of course it is. [laughing] We're talking about AI, you know my views on AI and my, my love of it, and everything it stands for.

Jan: Well.

Paul: Yes.

Jan: Can we change your mind do you think?

Paul: Well, let's find out if there are some proper uses for AI that are not just practical, but worthwhile and can make a real impact in the world of sustainability, shall we?

Jan: Yes, please.

Paul: Yes. Let's bring back a very old friend of the show, not that he's very old, it's just he was one of the first episodes, uh, that we spoke to him. Someone who is from the Department of Accounting and Finance here at Lancaster University, who previously spoke to us about his work on modern slavery reporting and regulation, and who is now back to explain how AI can be used in that.

Hello, Dr. Mahmoud Gad.

Mahmoud: Hi, Paul. Hi, Jan. Thanks for having me. I like this introduction because it's more realistic than the ones that you mentioned a few moments… [laughs]

Jan: [giggling] …yes, indeed, yeah.

Paul: So are you trying to tell me that you don't have real world relevance, or bring real world relevance to everything you touch?

Mahmoud: [laughs] Well, no. I try,

Paul: Yes. Intellectual rigour? Is that more, is that more accurate…?

Jan: …yeah…

Mahmoud: …yeah. Uh, yeah…

Jan: …that's more recognisable, yeah…

Paul: Yeah, yeah. When you're marking your child's homework and telling them what they've done wrong, that's the intellectual rigour side of it.

Mahmoud: Yeah. Probably.

Paul: It's very good to see you again, uh, Mahmoud, and to have you here bringing us up to date. 'Cause I guess this is almost like the next step from what we spoke about last time about modern slavery reporting, regulation and analysis.

Can you just remind us of, when it comes to modern slavery reports here in the UK, what are they, who makes them and what they're meant to contain?

Mahmoud: So, so the Modern Slavery Act was kind of enacted 2015 in the UK. And the main idea was asking UK companies with large business in the UK, 36 million pounds or more, to produce an annual statement on the actions that companies are making to detect, um, remediate and kind of mitigate modern slavery in their businesses and their supply chain.

And that's basically, kind of, the main requirement for companies. Companies have to talk about the risk assessments and talk about the structure of the business, their supply chain, and how they are dealing with, with modern slavery, uh, indicators and how they are remediating them.

Jan: When we last spoke, you had produced a report for the Financial Reporting Council to help them understand the quality of the modern slavery reports.

But that was about two and a half years ago. So what have you been up to in this area since then?

Mahmoud: I think three things have happened since then. Um, there was updating the guidance, so the Home Office issued, kind of, an updated, more detailed guidance in March 2025.

And I kind of talk, I can talk later about this, but we have also been working since the publication of our report that was done for the Financial Reporting Council and the Anti-Slavery Commissioner on building an anti-slavery intelligence platform. Uh, with the goal of kind of scaling the analysis and the, um, measuring the performance of modern slavery statement, uh, quality.

Um, and we have used this platform to analyse a large number of statements and, uh, some of the results, uh, might be uncomfortable as well. So we continue to find the results that are not necessarily, uh, pleasant.

Jan: Paul, I don't know if you know, but we are, we are both accounting scholars.

Paul: [sarcastically] Really?

Jan: Yes. [laughs]

Paul: Really, really? I, I, seem to remember the first time, no, I don't know if Mahmoud was the first guest you brought on with the, 'and here's another accountant' [inaudible] talk about accounting…

Jan: …sorry about that. So traditionally when we've been trying to understand, you know, the quality of reporting, it involves, you know, a for a, for a cohort of companies, selecting the reports and reading them manually.

So I wonder if, for our listeners, Mahmoud, if you could say something about that process and then the limitations of approaching reading reports in that way?

Mahmoud: So as you said, uh, the work was mostly done manually. So whether that's by academics, even actually the investment community when they started doing this analysis was mostly manual, and, um, also NGOs. So most of the work has, that was done on these statements was typically manual.

Kind of the biggest issue is that with, um, manual analysis, you are looking at, at most a few numbers of statements that you can analyse manually and typically, slightly, a snapshot.

So you focus on the, say, hundred statements in 2025. So you are, kind of, you are limited by the sample. You are limited by, kind of, the time series you can analyse.

And, uh, of course with us as humans trying to interpret things, we kind of tend to have our own views on is this really an evidence of, say, remediation? We can have kind of different views. And then that's why some people try to ask, say, uh, more than one review or to review a statement. And someone has to kind of create an agreed uh, scoring for a typical, uh, criteria.

So it was kind of, a, an intensive process, uh, before, uh, using AI.

Paul: Does it create issues if you don't have the same person reading it in the first place? 'Cause if they as an individual have their own individual tendencies, biases, and then they've read half the reports and a different person has read the other half of the reports, with different biases and personal preferences, et

cetera. Does that create issues, if it's only been done by people then, rather than computers?

Mahmoud: It creates issues, and that's why kind of some people try to find sort of a moderator that try, looks at those different people's kind of, uh, assessments and try to moderate them into an agreed score.

But again, still, also even the moderator view might be, has some subjectivity, uh, kind of embedded in it.

Jan: And I guess the best you get is a, an agreed subjectivity.

Mahmoud: Mm-hmm.

Paul: Which is almost like you get a, a medium for every, or a mean or, a…

Jan: …yeah…

Paul: …for every one of the reports that's done. And therefore, do you end up where maybe you should have more of an extreme either end, but you actually end up with everything being moderated by people having slightly different views?

Jan: I think it's, for me, when I've been involved in these processes, either as a moderator or one of the, the initial coders, it's not so much that you get a reversion around a mean, 'cause that wouldn't be a very good outcome, 'cause as you say, it sort of homogenises it.

But there'll be some element that when I read it, I read it in a really different way to somebody else…

Paul: …mmm…

Jan: …in which case, if you Identify that you then have to go back and recode everything. So there's, there's quite a few, you know, cycles in and around to be able to, to get there.

I think the biggest number I ever coded was about 120 or so with, um, my dear colleague, Carlos Larrinaga. At the end of it, he asked me that, that to never ask him to do that with me ever again, [laughs] 'cause…

Paul: …I know how he feels…

Jan: [laughing] …I'm really lousy at coding and, and it just took us so long.

Mahmoud: And the issue is that basically, if you think about it, I mean, um, what regulators typically have in mind when they ask companies to disclose information, is to allow for the market forces, people, uh, investors, community, academics, to discipline that disclosure.

And if we are really have this sort of subjectivity and ability to analyse only a small number of statements that kind of affects how we can discipline, uh, that disclosure.

Paul: You know, you, Carlos analysed about the number of papers with you that we've done podcast episodes?

Jan: Oh, that's really quite funny. [laughs]

Paul: That's when he walked away and said, never work with me again… [Jan laughs]

Jan: …no…

Paul: …just, just a comment…

Jan: …but there's, but there's also a really important thing that, that Mahmoud just said then as well, and that's the idea of, issue opacity.

So if you're regulating for what you should see in a modern slavery report, because everyone's situation's different and, and you're wanting to get an outcome but you can't specify all the in-between bits of that outcome, then the guidance itself is sort of kind of open as well.

So, so you've got open guidance on a report that people choose to write in a particular way given their own sensibilities. And then you've got a bunch of people, so you can kind of see how the threads of this particular cloth starts to fray.

[Mahmoud laughs]

Paul: It's almost like judging in gymnastics or ice skating…

Jan: …yes… [laughs]

Paul: …wherefore, that someone performs in a way that they think is the best performance, but the judges might have a view themselves. They'll have told you what you need to come out with at the end, but how you go about it might be different and…

Jan: …yeah…

Paul: …therefore you get marked down, and why you have TV commentators sometimes say, that was the best display I've ever seen. And then the marks come out and the person they're talking about is dead last. [Jan laughs]

Well, obviously not quite the best perfor… yeah, you've got people with different viewpoints at either end of the scale, which I guess is where we bring in our old friend the artificial intelligence, Mahmoud.

Mahmoud: [laughs] Yeah.

Paul: Why? [everyone laughs] It's just a general question whenever I think, about AI, why? But, yeah, particularly with this one, why did you develop the AI programme?

Mahmoud: So, building on what kind of you, uh, and Jan have been just mentioning, uh, it's basically if you really need to answer the question, is the regime working? You really need, need to look at the whole sample. You need to look at the population.

Say, all publicly listed companies over, say, five, six years and do this kind of analysis. And this, doing this by hand is really difficult. And there was no study, uh, was able to do this work. So the goal was basically not to replace human judgement, but to scale what kind of human experts do when they analyse these statements, uh, kind of at a large sample, uh, so that we can answer the question whether the regime is working, whether kind of interventions that happened say in March 2025 made a difference.

So looking at a few statements, uh, with subjectivity kind of involved, uh, sort of does not really help us to answer that question.

Jan: In this context, what is meant by an AI platform? 'Cause, AI's attached to sort of like a multitude of things. So what kind of thing is this platform?

Mahmoud: Yeah, so I'd say kind of a naive thinking would be basically, I am going to put modern slavery statement of company X into an AI model, ask it to score it. And that's a horrible thing to do. Why?

You will find that, uh, AI models are going to make a lot of assumptions on your behalf, okay? What criteria the model is going to use to analyse these statements. What is actually an evidence and what statement? And what is this company, what industry it belongs to?

All of this information the model is going to make, uh, an assumption on your behalf, and most likely you won't know what are the, what are the assumptions. So there's a black box where all of these decisions have been taken on your behalf and most likely will end up with, uh, an incorrect score.

You can put the statement over and over again, and you will always get very different results every time, because different assumption applied every time.

So, uh, kind of, the platform we are doing, basically think of it as AI is like an engine. And we built kind of a, a whole harness around it, kind of the core around this engine, that guides this engine to do the work in the way that we wanted it to do. So this is kind of the, uh, idea of the platform.

So we give the model, here is the model, we give it some tools. For example, in our case, we give the model access to, uh, Companies House information. So the model can go there, can access information about the company from Companies House, including the, uh, company number.

And this company number, the model is going to take and look at the modern slavery registry. So these are all facts. The model can't make this up. So the model is going to take this number, go to the modern slavery registry, crosscheck some information, and these are part of the criteria that we, kind of, use to analyse the statements.

So these are all facts. The model can't make assumptions on our behalf. Then we feed the model criteria. In our case, it's now 60 criteria. And for each criteria we give detailed guidance. And this kind of builds on, that we did some manual work before, so we kind of, okay, now these are, this is what could be an evidence on this item, this what should not be treated an evidence on this item.

So the model now has kind of complete, sort of, think of it as guardrail around how it should think. But also we, kind of, uh, enforce the model to, uh, output the, um, audit trail of all the decisions that's been made. So everything is transparent.

So if you as a company, or if you as an NGO, wanted to know why this score for this particular criteria, you can know, you can get the exact quotes from the statements that are used to give that particular score.

Paul: So we talked about bias in the human ones. How do you avoid any kind of bias when you've given it this criteria that, that it needs to analyse and how it needs to analyse it?

Mahmoud: So this is basically like, okay, can we trust, what, kind of, why should we trust these results?

Paul: Mm-hmm.

Mahmoud: So we, we, I mean, you can't completely kind of, uh, really eliminate some bias, but we did three things that sort of tries to assure people that we are, kind of, did something that's not completely, uh, irrelevant.

So the first thing, we are grateful for the CCLA investment management. Their experts have manually scored, uh, modern slavery statements, and they give us this data.

And what we did, we used, uh, our model to score the same statements, and we measured the accuracy of our model to get as close as to what those human experts, uh, can see in a statement as evidence.

And we, kind of, use traditional measures of accuracy. And, uh, one number I can mention now is something called recall rate. So this is like, um, out of all the things that, uh, the model showed pick up, that are picked by human, how much you are able to pick up. And we have 94% correlate, which is, kind of, very highly acceptable in these kind of, of models.

Um, also we have, as academics, we tend to, kind of, submit our papers for peer review. So the pipeline, kind of the AI pipeline we used, uh, that's kind of, is the kind of the backup for this model, how the model works. We submitted to, uh, one of the largest AI conferences in the world. It's called AI, um, international, uh, the International Conference for Artificial Intelligence. It's been reviewed, uh, accepted, uh, , for demonstration track where we are gonna demonstrate the pipeline for the audience in this, uh, uh, conference.

Second, third is basically everything is transparent, so you can get exactly why a given score is given for a specific company, including codes from the statements that are used to support this, uh, score.

Paul: So you say that it's got a 94% recall rate compared to what humans should see.

How much would the average human see? Because I know, obviously you'd want a human to pick up 100% of everything, but humans don't, so humans make errors too. Is therefore the error rate, you know, that 6% isn't there, is that similar to what, say, a human reading a report would have?

Mahmoud: I don't have specific numbers, but because, you know, we, if you ask a human to score it, then there might be some kind of error rate. If you ask another human to kind of say, okay, is this correct? The other human will have his own biases, then you can't really judge.

What we use, kind of, what, kind of, the norm here is to look for what an expert going to score a given statement, and then what our AI model, how accurate is our AI model relative towards those experts…

Paul: …mm-hmm…

Mahmoud: …uh, are going to score. So the idea's kind of really not to replace them, but can we scale what those experts are doing, uh, at kind of really a larger sample.

Jan: Yeah. And where does this happen? I mean, does it happen, um, see, I don't know what any of the model names are, really. So does it happen on…

Paul: …Naomi Campbell, um…

Jan: [laughs] …you're date…

Paul: …Kate Moss…

Jan: …you're dating yourself, there…

Paul: [laughing] …yeah, I know. My knowledge of supermodels nowadays is minimal, at best.

Jan: So, so there are obviously commercial models where you can ask them questions as, as Paul did for our, our introductions, et cetera.

Does this happen in a commercial model or does it happen in some sort of other bespoke space?

Mahmoud: So this happens in, in a commercial model. And uh, we kind of, uh, in our platform, we kind of are using all the analysis being done, uh, in the UK.

Uh, so we use something called data, UK Data residency. So all the statements are sent within a data centre in the UK. So the analysis, kind of, really done in

the backend is done within the UK, so it's not sent to somewhere that we don't know. So kind of, uh, uh, GDPR compliant. And it uses kind of the compute power of these models to do the analysis in the uh, backend.

We use, um, Google Cloud platform to sort of do the analysis in the backend.

Paul: Approximately how much time is saved then by using this as opposed to going through the old way of doing it? How many papers could be analysed in the time it took to analyse one by a human straight off?

Mahmoud: So a given statement could go from several hours by a human to less than five minutes by, uh, the AI model. So I can give you just, for example, in an analysis that we are currently working on, we analysed all publicly listed companies that are required to produce modern slavery statements from 2023 to 2025.

That was roughly, um, 400 companies, 850 modern slavery statements, uh, in a day. Uh, so imagine if you, kind of, ask humans to do this work, uh, how long it's gonna take.

Paul: And even if you then need people to check and audit, you'd expect people to check and audit if someone had done this manually anyway, so therefore...

Mahmoud: …yeah…

Paul: …the time has been saved, but you've still got the human element on the end of it…

Mahmoud: …of course…

Paul: …to make sure that it's been done properly.

Mahmoud: Again, also with AI, of course, we need some human in the loop. So with these 850, kind of, statement analysis, you need some intervention at some point to make sure all is going in the right directions.

Jan: and I suppose what's really exciting about that is that most of our studies of corporate reporting have been on samples. Either samples of a subset of companies, or a subset of times, and often they are prepared not using the same research instrument.

So you might have results about a certain piece of reporting, but it's all varied 'cause everyone's done their own sort of study. Whereas it, is it unusual now we can talk about actually researching whole populations? Everyone's.

Mahmoud: Exactly. So this is something that I guess what's gonna come now with the scalability that AI is producing. Again, assuming that we have built the right car around the engine, uh, we can really do, uh, a lot of work in that sense.

And I can, I can tell you the models have kind of moved long way from, say, late 2022 or early 2023, because when you analyse, started this project was really with, say, the early models, uh, that were released in 2023.

And these models were hallucinating as much as you can, you can imagine [Paul laughs] how it can fake everything. And, basically it had a very kind of short context window. So, um, you can put a statement to a model or a large text to a model. So we used to use something called retrieval-augmented generation, early stage.

So what's the retrieval-augmented generation, is basically you are going to, instead of feeding the whole statement to a model and asking it, what are the, if the company has remediated, uh, the identified cases of modern slavery, you are going to start with a semantic search process.

So the semantic search process is gonna take your question, look for all the, uh, similar discussion within the statement, and you're going to summarise these, kind of, say, 20 sentences and pass *only* these 20 sentences to the model, not the whole statement.

Because the model had a very short context window, it cannot really process more than say, 16,000 tokens, I guess, which is like less than, uh, which about 15,000 words, more or less. So now context window can go to about a million token. Some of those can really read a full modern slavery statement. You can really force the model to follow your instructions so much better than before.

So things are improving, and this will take us into this process where we can analyse the whole population with some uncertainty, but we can really control this level of uncertainty.

Paul: Imagine if you'd had a similar thing when you were doing all your analysis in those dusty Scottish archives looking at, [Jan laughs] you know, Highland clearances, and what have you, in the 17th, 13th centuries…

Jan: …yeah. I mean, it is a game changer in that respect, which is why, because I've, you know, sort of start off in general, you know, both ignorant and hostile to, um, AI and, and use in these kind of applications.

But I've, I've, becoming more informed by conversations with Mahmoud and with others, and think, ah, there might be something here that can really increase the, the reach and the rigour of work.

Paul: So if it is to be useful, who's going to use it? Who do you see being able to use this database, this, um, platform and, and how?

Mahmoud: I would say at least, uh, regulators who wanted to have a macro view on how things are happening, in terms of modern slavery, uh, reporting. Uh, is the system working or not? I think these, kind of, whole population analysis will give some indication.

I would say also companies themselves, because what our platform kind of really kind of produces is basically a score for a company. And then how you score compares to your competitors, like, industry peers, but also similar size firm.

So as a company, say, if you're a manufacturing company, you will get your score. And how your score looks like relative to all other manufacturing firms. But also if you're a FTSE 100 firms, how your score compares relative to those FTSE 100 firms.

And then you'll get some feedback on, okay, what are things that looks good? What are things that, kind of, you can work on and how can you improve these things? So I guess kind of companies, especially, kind of, mid-market firms that are, kind of, invisible to scrutiny by, kind of, investors' benchmarks.

I would imagine those companies will find, uh, some value in it, because they don't have, like, large compliance teams that can look at the statements and, and really, uh, work this out. So I guess these are part of the beneficiaries of, of the tool.

Paul: Had a nasty flashback then to benchmarking and your, [Jan laughs] your seafood, uh, benchmark table, where the reports were done and what companies could do based upon that. By comparing themselves to other ones, yes.

Jan: But I think the other really interesting thing you said is about these mid-market firms. So that might be a phrase that's unfamiliar to our listeners.

So maybe say who they are. 'Cause, uh, 'cause they are, if you like, the bulk of the activity, and the people who have, as you say, the less resource. But, but what is a mid-market firm?

Mahmoud: Well, it can be defined in different ways. So in our case, we are looking at firms that, say, below FTSE 350. So if you're not in the top 350 firms in stock market, you are a mid-market firm. But you can define it in very different ways, uh, by assets, by total sales, by number of employees.

But we argue that these are the companies that don't have, that are not kind of in the spotlight, but also don't have enough, kind of, teams that kind of do this work, uh, for, for, uh, compliance work on, on modern slavery statement.

Paul: So bigger than like your micro businesses, your really small SMEs, but not at the top scale…

Mahmoud: …yeah…

Paul: So those ones, like you say, right in the middle of it all.

Jan: But then it becomes more of a learning tool for them. So, well it might be an evidence tool for, for the reporting council, it's a, you know, help me do this better tool for companies.

Paul: What's the reaction been then from the reporting council? 'Cause obviously you worked from, with them, at the very start of this process.

Have they been along for the ride with you on using AI and seeing how it might be worked with, with them?

Mahmoud: So we have just finished kind of, we have just kind of, the tool was out only kind of recently and we have just finished the analysis and we are kind of in the process of writing this up.

And the goal, hopefully we're gonna send them the results of our analysis, uh, to bring attention, uh, to the tool on, or in what we, what can be done in this space.

Jan: So what's next for you in this, this exciting journey?

Paul: So you've gone from manual reporting to AI reporting to, what comes, yeah, what comes after AI reporting? [Jan laughs]

Quantum computer reporting…

Jan: …I was, so, I wasn't quite on that trajectory…

Paul: …oh, I get you…

Jan: …but, but I could see… So what next for this research agenda, but also for your work in this area?

Mahmoud: I do think for the research agenda is to try to publish this report, this analysis of 850 firms, because kind of some of the results there might be informative.

So I'll give you just one example. Uh, we followed companies, uh, in the due diligence chain. And you see basically, among these companies with, uh, that says we do, say we do a rigorous risk assessment. So these are 67% of firms in the, uh, in the, um, in our analysis, only 67%, uh, says that we do a rigorous assessment.

And among these firms, uh, 40% only provide information about the audit protocol. And then as you move down, grievance, uh, whistle blowing and identification, we are talking about only 9% of firms identify actual modern slavery. Um, and then 4% remediate, and then 2% provide kind of satisfactory remediation.

So you can see, kind of, this is like a funnel. And it keeps collapsing as you move down the, uh, due diligence chain, especially when the evidence is really needed. So we have a system here and you say, okay, I'm doing rigorous assessment, but you're not detecting. And given that modern slavery is really endemic, it's really strange that you are doing all of this risk assessment and you're, you're not finding modern slavery.

And what we found this also is that, basically companies with low scores in terms of due diligence in 2023, continue to have low scores, uh, on kind of 2024, 2025. So this is one of the things that kind of poor reporters continued to be poor reporters.

On average companies, kind of, seems to be marginally improving. If you look at the whole market there is kind of marginally improvement in terms of the scores. But what's really interesting is that 95% of firms that say we find indicators provide really detailed mitigation steps.

So, kind of, the bottlenecks here is for companies to say, we identified cases of modern slavery. So once companies identify this, kind of, we think this would

lead to detailed mitigation steps, which is kind of, uh, the whole goal of, uh, of the process.

Paul: Well, Mahmoud, thank you so much. You've brought us up to date and seen that next step from the work that we talked about almost three years ago now, and it does seem like there might actually be a use for AI that's reasonably useful…

Jan: …well done, Mahmoud, well done…

Paul: …that's about as high praise [Jan laughs] as you're gonna get from me on artificial intelligence.

Thank you, Mahmoud.

Mahmoud: Thanks so much. Thank you. Thanks, Jan. Thanks, Paul.

[Theme music]

Paul: In just three years, you can see how much Mahmoud's work has moved on, and how the use of AI has allowed him to build upon what he already had, and take it to another level, and get so much more use out of it.

Jan: Absolutely. And it sounds like that's both the models moving on, but then the approaches to using the models as well.

Paul: Mm-hmm.

Jan: And I, I really liked his analogy of, of an engine. And it's what you've built around it that makes it, you know, robust.

Paul: And I liked what you said as well. 2023, it was producing so many essentially mistruths and fictions and um, they've worked on it over three years.

And if you imagine with the AI that we're all using in our lives, how much generalities there is there, whereas this is on a very narrow subject, and it's taken them three years of concentrated work on this narrow area to get it to a stage where those fabrications and are no longer there, and everything's being produced to that great level of accuracy, 94% of what should be being picked up is being picked up.

Brilliant. That demonstrates the great task there must be with AI on the global scale of making sure that it's not making up anything.

Jan: Yeah. And I think it's that differentiation between like a general, you know, ask and see what you get, and actually having it under these guardrails and, and supervising it very closely…

Paul: …yes…

Jan: …which I thought was amazing,

I really liked some of the findings, though, they had at the end of this process as well. And one of the concepts I found really engaging is the idea of a collapsing funnel.

Because I've seen collapsing funnels in all sorts of other areas of work as well. So looking at, um, Tim Lamont's work that looked at the restoration practices of companies, and what they say about it, that had this collapsing, um, process going on as well.

But likewise, looking at, um, biodiversity disclosure by ocean companies, by indicators disclosure around the TNFD. So, collapsing funnels is something I think, you know, for a select number of our, our listeners will go, oh yeah, let's gonna use that phrase, 'cause it captures it.

The other thing that I really liked is that he talked about bottlenecks. So if you are unable to disclose this particular piece of information, either 'cause you don't know what's going on, or, or you decide not to say anything about it, then everything else sort of wanes off.

So that bottleneck idea, I think, is a really important one as well.

Paul: Yeah, I know what you're saying about the, the collapsing funnel, some of the numbers he was giving out there where you've got from 67% of firms who were the reporting bit, doing it OK. Down to, there's only 2%, I think was the last number he gave where he'd gone through various things, such as um, whistleblowing, and um, stuff like this, and then mediation and then satisfactory mediation was the last one, on 2%...

Jan: …yeah…

Paul: …yeah, that's showing that whilst people are at the top level, the people who've gone all the way through the process to get, uh, just…

Jan: …yeah…

Paul: …such a small amount, but still lots of work to do.

And he did. You know, calm some of my worries about all the AI biases in there. They, they've obviously put an awful lot of thought and consideration into how to account for the biases by having so many people involved who are experts in the field, getting it fed in.

It's not just that one voice that's putting all the information in and telling the AI what to do, and therefore the AI will produce things with that person's voice. They've got lots of expertise in there, and I was really happy to see that.

And it's good to know that if they're able to use this to get firms to climb up the funnel. I don't know if that's, that's the way…

Jan: [laughs] …yeah…

Paul: …uh, rather than sinking through the plug hole, um, then even better…

Jan: …yeah, yeah…

Paul: …if they're able to say, look, this is what our report says to you, you could do this and you will climb up a bit…

Jan: …yeah, yeah…

Paul: …and then that's what we can do.

Jan: Yeah. No, I thought, well, that, that potential to unlock better reporting and we, we believe, 'cause there is some evidence, but, but more evidence is always welcome, that good quality reporting, um, often, uh, replicates in good quality practice.

Not, not straightforwardly, but there is a connection there that could be drawn out.

Paul: Well, we can see the end of *our* funnel because next week is our last episode…

Jan: …heavens…

Paul: …of this season…

Jan: …wow….

Paul: …so, yes…

Jan: …who are we going to talk to…? [laughs]

Paul: …well, I, I was gonna talk to this strange lady I met once on a train. Oh, no, that wasn't me. [Jan continues laughing]

Uh, it's just me and you, Jan…

Jan: …well…

Paul: …we're, we're gonna wrap up, discuss all the various topics that have appeared a long way through lots of episodes.

Jan: Oh, that's, that's tremendous. And how many episodes have we done, more or less?

Paul: Um, this will be episode 43 of this season…

Jan: …heavens! Just like that…

Paul: …yes, yeah, 43 weeks of me and you, and [Jan laughs] various victims who've sat in the chair next to us....

Jan: …well, I look, that, that sounds great.

Paul: Yes, you can be the victim next week.

Until then, thank you very much for listening. It's goodbye from me, Paul Turner.

Jan: And me, Jan Bebbington.

[Theme music]