



LSEG STREETEVENTS
EDITED TRANSCRIPT
TRUSTPILOT GROUP PLC CAPITAL MARKETS EVENT

EVENT DATE/TIME: May 06, 2026 / 1:00PM UTC



An LSEG Business



CORPORATE PARTICIPANTS

- **Adrian Blair** Trustpilot Group PLC - Chief Executive Officer, Executive Director
- **Shazadi Stinton** Trustpilot Group PLC - Chief Trust Officer
- **Maj Santhakumar** Trustpilot Group PLC - Senior Director Trust Operations
- **Thomas Vermeulen** Trustpilot Group PLC - Director Fraud & Investigations
- **Sona Pakhchanyan** Trustpilot Group PLC - Senior Applied AI Manager
- **Domonique Rai-Varming** Trustpilot Group PLC - Vice President - Trust & Legal
- **Hanno Damm** Trustpilot Group PLC - Chief Financial Officer, Executive Director
- **Ciaran Dynes** Trustpilot Group PLC - Chief Product Officer

CONFERENCE CALL PARTICIPANTS

- **Hai Huynh** UBS AG - Analyst
- **Timothy Ramskill** Bank of America - Analyst
- **Sean Kealy** Panmure Liberum Capital Ltd - Analyst
- **Ross Broadfoot** Rbc Capital Markets - Analyst
- **Dan Ridsdale** Edison Investment Research Ltd - Analyst
- **Thomas Brown** Premier Miton - Analyst

PRESENTATION

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

All right. Good afternoon, everyone, and a warm welcome to all of you and those joining by webcast. Thanks for being here for what I hope will be an interesting afternoon for everyone. We have a packed agenda during which our Chief Trust Officer, Shazadi Stinton and members of our Trust team will bring to life the way we make Trustpilot the world's most trusted open customer feedback platform.

Our CFO, Hanno Damm; and Chief Product Officer, Ciaran Dynes, will talk about the benefits of being an open platform in the age of AI. And we'll then move on to a Q&A session in the room. Those of you online will be able to submit questions through Spark Live, which we can answer in the room. We then have some breakout sessions to dive deeper into a few areas with the opportunity to ask further questions. Finally, we hope those of you with us in person will be able to join us for a few drinks at the end. So let's kick off.

Trust is and always has been the most important thing in business. Whatever industry you're in, trust is what makes people want to work for you, invest in you and be your customers. The fundamental ways in which businesses build trust are also timeless, be transparent with information, be credible on by delivering on your promises and treat people with respect as humans. But in the age of AI, it's easier than ever to look and sound credible. So deciding which businesses to trust becomes more challenging.

Trustpilot has a critical role to play as a governance space in between the chaos of social media and the bias of information sources controlled by businesses. As large language models amplify the feedback on Trustpilot, our platform is becoming the critical trust signal for the age of AI. In essence, Trustpilot products collect customer feedback and turn it into influence over the actions of people and businesses. Because the feedback has influence, people are more motivated to write it and businesses want to engage with it.

The more AI advances, the more influence that feedback has. For almost 20 years now, we've continually invested in the know-how and technology needed for this work. From cutting-edge AI to specialist human expertise, we will do whatever it takes to keep the platform trusted as the world around us changes. While today is an educational event, it's vital to understand this. Trust underpins all of our performance. The evidence you will see this afternoon is what ultimately gives us the confidence in the targets we set out to the market back in March, including reaching an adjusted EBITDA margin of 30% by 2030.

Today, you're going to hear a lot about the work we do to maintain trust in the platform. You'll be wondering how do we measure the success of all this activity? The ultimate yardstick rather than being a single metric is the simple fact that with complete freedom of choice, millions of people and businesses find the feedback on Trustpilot useful. CEOs get their teams to read and act on the feedback because it's genuine.

People turn to Trustpilot again and again when deciding where to spend their money because they actually learn something useful from it. The platform, as you can see here, has grown in scale and influenced more than ever since the rise of large language models. In 2025, our users submitted more new reviews to Trustpilot than in the first 12 years of the company combined. Our unique data set expanded to 361 million active reviews, representing a 20% year-on-year increase.

The volume of activity on the platform isn't just a sign of how much Trustpilot is trusted. It also gives us the data to sustain that advantage. As we'll see today, the volume of feedback on Trustpilot is critical to our ability to identify and remove content that violates our guidelines and so is itself a valuable competitive moat. Today, we're going to take you through the systems, capabilities and discipline that make Trustpilot the trust layer of commerce in the age of AI.

To do that, I'd like now to hand over to our Chief Trust Officer, Shazadi Stinton, and her team. Over to you.

Shazadi Stinton Trustpilot Group PLC - Chief Trust Officer

Thanks, Adrian. Good afternoon, everyone, and thank you, Adrian, for the introduction. As Adrian mentioned, I'm Shazadi, and I'm the Chief Trust Officer here at Trustpilot. Today, I'm going to focus on how we build trust at scale. I lead the team responsible for trust, transparency and platform integrity.

Our role is to ensure the platform consistently delivers authentic, reliable information at scale. This is supported by a trust infrastructure organization of over 200 specialists across policy, enforcement and technology. At its core, our challenge is simple: how do you maintain trust in a global platform operating in real time and at high volume?

Today, I'll cover how trust is built, how it scales and why it's increasingly critical to decision-making. But first, let's define what we mean by trust. We see trust is built on three things: transparency, credibility and humanity. Transparency starts with visibility. People trust what they can see and what they can see in full, not just the positive, but the negative and the unresolved. The moment information feels curated or incomplete, trust doesn't build its stores. For us, transparency means exposure to reality, not a polished version of it.

Next is credibility. Credibility is what makes that visibility usable. It's not enough for information to be visible. It has to hold up. That means consistency over time, signals that are difficult to manipulate and systems that behave in ways people recognize as fair. Credibility is what turns something people can see into something they're willing to act on.

And then there's humanity. Trust is ultimately a judgment about intent. Are businesses listening? Are they responding? Do they acknowledge when something goes wrong? Humanity is what turns a transaction into a relationship. For us, they are nonnegotiable, get them wrong and no level of product sophistication can compensate, but get them right and trust compounds over time.

So let me now turn to Trustpilot itself and why it matters more today than ever before. The Trustpilot, we're an open independent platform where consumers and businesses come together around one thing: genuine customer feedback. Since 2007, we've built one of the largest review data sets, over 361 million reviews, with more than 190,000 added every day. But scale is only meaningful because of what it enables. We help consumers make better decisions at key moments. And we help businesses understand feedback and improve their performance. Our ambition is simple: to become the universal symbol of trust.

To make trust real and scalable, we operate with five core principles, which underpin everything we do in trust. So we are neutral, our systems do not favor consumers or businesses. We're open, anyone can share a genuine experience and businesses can respond. We're fair, the same rules apply to all users regardless of who pays. We're transparent, we explain what we do and why. We publish clear guidelines on how we operate, moderate content and calculate the trust score, and we're relevant.

We continuously improve the experience for both consumers and businesses to ensure the platform remains useful at scale. These principles are embedded in how we design systems and enforce policies at scale. So let me ground this in how the system actually works. First, every business is held to the same standards. There are no preferential rules for paying customers. Second, moderation is based on evidence and consistent rules. It is not influenced by commercial relationships. And third, ratings are earned through customer experience. Businesses that actively engage with customer feedback and use it to improve their service tend to perform better over time.

They build stronger customer relationships, increased retention and are better positioned to adapt to changing expectations. A trust score is, therefore, not just a rating. It is a reflection of real customer experiences and how businesses respond to them.

So at the center of our platform is a really simple loop. A customer has an experience, they leave some review -- they leave a review and that feedback is published transparently. Businesses that respond and improve typically deliver better experiences over time. Those improvements then show up in future reviews. Trust, therefore, builds through behavior. So we provide paid tools that help customers collect and act on feedback, but those tools do not influence what is published on the platform. The system only works because users trust that reviews reflect real unmanipulated experiences.

Now trust is being rebuilt in the age of AI. What was once a moderation challenge at scale is now an adversarial environment where organized bad actors use AI to generate fake reviews, synthetic identities and coordinated the campaigns instantly and at scale. The cost of this section has collapsed. That changes the game. Winning platforms won't just react faster. They'll use AI to stay ahead, detecting patterns earlier, acting in real time and improving decisions continuously.

At Trustpilot, we've built for that reality. So we operate trust infrastructure not just as a review platform, continually assessing authenticity, detecting abuse and enforcing standards in real time. This is powered by our proprietary AI trained on years of unique data designed to both identify bad actors and improve overall platform integrity. So we balance three priorities: protecting consumers, ensuring fairness for businesses and maintaining transparency in how we act.

Today, we'll show you how that works in practice. Our safeguarding ecosystem and how we apply our policies at scale, how we address misuse by businesses, how we use AI to detect and prevent fraud and how we stay ahead of a rapidly evolving regulatory landscape. We'll close this session with a conversation between Hanno and Ciaran on why openness wins in AI. So our belief is simple. Openness will define trust in an AI-driven world.

Finally, we will move to three breakout sessions covering fake review sellers, business verification and community flagging. So the key message I'll leave you with is this: trust at scale doesn't happen by accident. It is engineered, tested and continuously improved. And as deception gets cheaper, the ability to defend and strengthen trust becomes a critical competitive advantage.

So I'm now delighted to hand over to Maj, Sona, Thomas and Dominique to take you through the details. Thank you.

Maj Santhakumar Trustpilot Group PLC - Senior Director Trust Operations

Good afternoon, everyone. I'm Maj Santhakumar, Senior Director of Trust Operations. Today, I, along with my colleagues, Thomas, Sona, who unfortunately can't be here today, but will be joining in via recording, and Dominique are going to speak and give a sneak peek under the hood of how we protect trust in the platform.

Our vision is to be the universal symbol of trust. And over the next 30 minutes or so, we will show you the infrastructure that makes that possible at scale. Because we are an open platform, anyone can write a review, whether invited or unprompted. With that openness comes responsibility. All users must follow our guidelines, whether they are reviewers or businesses.

To enforce these guidelines, every review on our platform is moderated through our three pillars, also known as our safeguarding ecosystem: our technology, our experts and our community. Our technology helps us detect fake reviews, apply our policies consistently at scale and identify potential misuse by businesses. Our experts support this by adding context, nuance and judgment where automation alone isn't enough. And our community plays a vital role by flagging content based on their own experience and knowledge. Ultimately, technology is the foundation that allows us to protect trust at scale. But it's a combination of technology, people and community that makes our platform resilient.

Now, I want to lift the lid on what happens behind the scenes with our technology when a review is submitted. Every review enters a 2-hour posting delay before it appears on the platform. This gives our automated systems time to act and prevent bad actors from learning and adapting to our detection methods. During this window, reviews go through two key layers of checks, fabrication and guideline enforcement.

At Trustpilot, accuracy in our detection systems isn't just a technical goal. It's fundamental to maintaining trust across our platform. Some of the best published academic results on comparable data demonstrate around 1 in 100 false positives at fake review detection tasks. While that may be acceptable in a controlled research environment, applying that standard at Trustpilot scale would have serious real-world consequences.

With millions of reviews flowing through our platform, even a 1% false positive rate could result in hundreds of thousands of authentic consumer voices being mistakenly removed. That's why we hold ourselves to a significantly higher standard. Every genuine review represents a real customer experience and removing it incorrectly risks undermining confidence in our platform. We detect, fabricate reviews using behavioral patterns, such as businesses reviewing themselves or coordinated review seller activity. These models continuously improve through human validation, our own research insights and feedback from user appeals.

At the same time, we operationalize policy at scale using AI. This includes detecting advertising or promotional content, highly sensitive personal data and harmful or illegal material, areas where consistency of decision-making and speed are critical. Making all of this possible, we have humans in the loop. Our internal experts conduct proactive investigations and validate edge cases, including training and refining our AI and LLM-based decisions. Together, this combination of automation and expert oversight is what allows us to protect trust at scale.

Let's show some examples of our technology in action. These are real reviews caught by our automated technology during the posting delay. AI enables us to catch this content before any genuine user can see them. I'm sure everyone in the room will agree, content like this has no place on Trustpilot. And we work hard to remove this before anyone can see it.

I'm going to pass on to Thomas, who will look through -- who will go through how we tackle business misuse on our platform.

Thomas Vermeulen Trustpilot Group PLC - Director Fraud & Investigations

Good. So I'm Thomas Vermeulen -- and the name maybe sounds not really English, I'm not. I'm from the Netherlands. So there is a bit of an accent, I'm sorry for that one. But I'm going to talk about business misuse today because Maj and others have talked about fake reviews. But actually, what we see when it comes to platform, it goes beyond the fake reviews. There is a whole work when it comes to people that are misusing the platform, for example, the way of inviting.

And what I want you to do is to show you actually what that is. So business misuse, otherwise that fraud, what is fraud. Fraud comes in various ways. And that's what I'm trying to visualize here. We've got review sellers, very visible when we have our protection against them. Fake and illegal businesses, for example drug dealers, illegal medicines that are being sold, scam operations, impersonating Trustpilot. Those are fraud types that you will see in other businesses as well.

What is specific to Trustpilot is business misuse of our platform. And I touched upon it very slightly, but I would love to dive into it a little bit more. So it's not only about the fake reviews. It is about how are you collecting your reviews. Are you doing that in a fair neutral and unbiased way. The information that you have on your profile page, is that actually right? Are you presenting yourself with the right information? So business misuse is covering a whole area.

And actually, this is never going to visualize at all because fraud is constantly changing. What fraud is today is not what it is tomorrow. It's evolving. It's getting new kinds of fraud, new techniques. So it's our job to continuously follow up, detect and protect. And when it comes to protecting, we need to have a very clear policy and enforcement. My background is in law enforcement, so I feel very comfortable to set clear guidelines on what is allowed and what is not allowed. And that's also what we have done with Trustpilot.

We have created very clear policies. And those are the ones that are reflected as guidelines for businesses and they're in our action we take policy. And actually if you breach those ones, that has consequences. And that's what you see here on the side. So we start with education. Because, yes, there are people that are doing things that are wrong, but they simply don't know. That's where education is really important.

But the bad actors are the ones that will continue and those are the ones that do it deliberately. So that's where we -- if we continue to detect the misuse, we will move on to the next step. That's a warning. Then, again, if that continues, we go on to the legal notice

and then it becomes serious. That is your last chance because if you then continue, we will have to take more clear and strong actions. And that means that we will terminate the contract. Whether you are a paying customer or not, it doesn't matter.

My team is not connected to commercial. So I've got the joy to say, you're doing it wrong. That's out. And that means that we terminate the contract. We limit the access to the platform. We place a clear consumer warning on the profile page. We hide the TrustScore because actually the TrustScore, that's the tool that people use to gain trust. Why do people use Trustpilot? It's to gain a level of trust from consumers. So if we would still show the TrustScore, then they still get the benefit. That's why we hide it.

Another very strong one into that one and especially in the age where we are right now is search data. So once we go this far, once we have detected, once we have terminated, we also stop the data sharing. So people will not be, through our data, visible on Google.

And then finally, there is an opportunity for us to go even beyond and to take legal action. We do collaborate with authorities. We report people. And we have even taken legal action against, for example, review sellers. Specific to the review sellers, we will talk about that in the breakout.

Now let's make it a little bit visible. So I have shown you what kind of actions we take, but how does that look then? This is -- once the business has continued the misuse, this is the final stage. This is where we have transparent and clear told you why we have terminated this business and what they have done wrong.

And actually, if we then look into actions because we do need to talk about actions to make it clear that we really are serious about this. On screen, you see numbers, ineligible businesses, businesses that we believe shouldn't be on our platform, either because they don't meet our core values or because they are illegal and fake businesses. What you all see is consumer warnings. We see what we have done off platform.

Let's dive in a little bit. So we have removed close to 12,000 businesses over platform straight because we believe they shouldn't be on our platform, illegal businesses, for example. We have placed 10,338 warnings to inform our consumers about there is a risk on this business or, hey, this business has misused our platform. Fake reviews, 7.8 million reviews we removed in 2025.

And to show you that we are really serious about this. And this is when my passion comes out, right? We are making this in a much more efficient and more scale year-on-year. And that's what you see in the increases: 184% increase on eligible businesses, 20% increase on where we are actively informing our consumers about the misuse.

Now in the review session, we will go a bit more into detail, so I'm going to swipe on to the next piece. And this is where it's going to be an interesting one because this is where you've got to play a role. We're going to play a quiz, which is called is it real or is it fake? What are we going to do is you're going to be part of my team today. You are all going to be a fraud investigator or you're going to participate as if you are our software.

And what you're going to see is you're going to get two reviews. And we start with the content of the review. And along the line, I'm going to show you more insights and more data. And every time I ask you to think, is the first one real or is the second one real, how do we do that? Slido. I have to have use slido. So I'm a bit new to that. But if you scan the QR code, then it should allow you to get in this one. The screen that allows you to say or to answer every time which is the fake review. All right. Here it comes.

So read the reviews carefully, two reviews, both reviews is about a headset, headphones and so on. And now the question is, based on what you're seeing here, you can use the star rating, you can use the content, you can use the name. Which review is fake? Is Review 1 fake, then select Review 1. Or if you do believe that Review 2 is fake, then select Review 2.

And in the meantime, I keep the pause. I keep an eye on the results. I might get you to see if the results are there. So 65%, it's moving a little bit, believes Review 2 is fake.

Now I'm not going to reveal it now, but I'm going to give you a hint. And I'm going to talk a little about this because if I look at this review, I see spelling mistakes. Now spelling mistakes doesn't necessarily tell you whether it's a fake review. Actually, if I write English, I might put in a small mistake sometime. If I write Dutch, I might be better. Review 1, that may be the one that's created by AI.

Sometimes AI does give you an indication, right? If the review is way too perfectly written, way too many punctuations, you could say, well, that's an AI-created review. But even then, isn't AI-written review straightaway fake? I use Gemini and ChatGPT to write my notes and stuff like that. So I could use it. Luckily, if it helps me. I definitely have a better way this one today.

Luckily, as in the fraud team and our detection models, we have much more data. So what you got right here is an extra set of signals. You've got the name. You've got the location from where the review is written and whether or not the reviewer has a profile

picture. Again, the question for you is, which is the fake review? Is Review 1 fake, then select Review 1. Is Review 2 fake based on the new data, then select Review 2.

All right. I think we're already there. 75% still believes that Review 2 is fake. Now again, I'm not going to reveal it now, but I do want to say something. Because what I see here is Sofia M, almost sounds like a criminal, Emily Brown and I see a profile picture. Now, it could be that this person, Sofia M, is actually from Spain and this is her favorite spot. At Trustpilot, you are free to select what kind of profile picture you want. Of course, it can't be harmful and illegal, but there is some freedom. So it could be that you selected this because this is your favorite place.

So let's move on. Third set of data points. What you see here is the reviewer history and some additional insights about the location. Again the question for you, Review 1 or Review 2, which one is fake? If you believe it's Review 1, then click on Review 1. If you believe it's Review 2, click on Review 2.

History tells a lot. But now actually what we see is a little sneak peek into our detection ability. What you see onscreen here is that Reviewer 1 is connected to five other reviewers. And how do we know that? It's because we use fingerprints. So the device settings, the device fingerprints, allows us to say Reviewer 1 is actually connected to Reviewer 2, 3, 4, 5. It could be that they all used the same device. I won't go too much into detail about that. What we also see is that, hey, they also reviewed -- all 5 reviewed the same business. While fake -- well, Review 2 is actually just one user that is connected to another user.

But if we look at what e-mail address that is, and we do see that they have the same family. Now does this change your mind? So again, which review is fake? Is it Review 1? Or is it Review 2? And what we see right now is that we went from 75% believes it was Review 2 that was fake, we see it hop to the other side. We actually see that close to 80%, 77%, 78% says that Review 1 is fake.

All right. Let's give a last shot at it. We've got to push quite hard on this one. So another piece, another data signal. What we know based on our detection and the way we look at data is that actually the Reviewer 1 was written from a disposable e-mail. That means the disposable e-mail is just a onetime use e-mail address. It doesn't really exist. And if you look at how it is created, it's totally random where actually Review 2 used a widely recognized hotmail.com and a very normal e-mail address.

Last time, which one is the fake? Fake is Review 1. And what I have tried to show you in a very quickly -- in a very quick way is actually what I'm trying to show you on the next slide. When people use our platform, they might see around five data points. They might see the name. They might see the content. They might see the location of the review. Those are the data points that you have if you look at Trustpilot, as anyone look at Trustpilot. And what we have actually seen in this trial that we did is that with just those data points, it is super difficult to identify which one was fake.

Actually, we all believe that fake -- that Review 2 was fake. This is what Trustpilot looks at. We look at device settings, connectivity, location, steps, e-mail addresses. It even goes as far as that piece of battery on your iPhone. It's that to say is for all reviews they present, that's a very typical signal for us to say, why would there be so many reviews when it is exactly the same amount of battery.

So what I'm trying to say is spotting fake reviews is really difficult. If you just do it with the top of the iceberg, the review content and maybe those other patterns that I've mentioned before, you won't succeed. But that's why we use all the data points that are sitting below the surface. That's why we also use the technology to do this at scale. That is how we are able to detect the fake reviews. And that should be my part.

Then I'm going to hand over to Sona. Well, Sona is not here, but we made a good recording, I believe. And she's going to present more about our AI and our automated detection.

Sona Pakhchanyan Trustpilot Group PLC - Senior Applied AI Manager

Good afternoon. My name is Sona, and I'd like to talk to you about how we use AI in our fraud detection systems. As Maj said earlier, high precision is a necessary requirement of our fraud detection systems, as we aim to operate at an extremely low rate of false positive detections. This requires building custom solutions tailored to our data, and it requires having a research pipeline that continuously delivers innovation to keep up with the evolution of fraud.

Fraud isn't a static target. It evolves. There are four key fraud behaviors that affect how we think about technology. First, the way fraud looks changes from week to week. So our models need to retrain weekly to keep pace. This runs as an automated process with a decision gate that releases new models only if they perform better on new data.

Second, fraud adapts to our defenses. As soon as we close one attack vector, the fraudsters start engineering around it. We compensate with rapid tactical rule releases, surgical and fast to deploy, sometimes as fast as the same day we spot the fraud

pattern.

Third, fraud targets the weakest signal. Whichever area of detection hasn't been reinforced recently becomes the new attack vector. So each fraud test requires its own dedicated defense. But none of these defenses stand alone. They feed each other, share signals and learn from each other. The whole system evolves as a single defense net. That's our way to stay ahead of the adversary.

Fourth, fraud spread through networks. Bad actors share tactics, tooling and playbooks. So we built early threat detection that flags emerging coordinated patterns before they scale. What you're seeing here on the left is a real time lapse of fraud network evolution from the past year. From week to week, you can see the fraud pattern rapidly developing before the defense kicks in to contain it. Every new review decomposes into multiple feature types: text content, user history, review timeline, device AD are just some examples of the types of data we assess for each review.

Those features feed our processors, content analyzers, graph builders, metadata extractors, which then feed various types of classifiers running in parallel. One example is our graph neural network, a type of deep learning model known as a graph attention network, where each node importance is calculated using a weighting known as attention, the same key architectural component as the GPT models.

These are great at learning to recognize a wide range of suspicious patterns expressed in relationships between businesses and consumers. In addition, we have models that we refer to as big impulsive trees that use decision trees to handle complex feature interactions and can, therefore, extract meaning from a wide variety of individual fraud signals, combining them to get high confidence decisions.

Other models we use include large language models, sequence models and hand-authored SQL rules. Each classified cast an independent vote. If any of them flags a review, it escalates. A trust check applies a last chance override for reviewers with strong legitimate history and the outcome published or filtered. An important part of the system that rarely gets the spotlight is what happens after the outcome. Every decision feeds into a pool for audit samples. We check those samples regularly for false positives and false negatives. Then many of these reviews feed into the retraining pool. The models you saw at the top of the screen learn from the system's mistakes and successes every cycle.

To summarize, there are 7 stages in our fraud detection cycle. Platform monitoring at the top triggers alarms when something looks off. Experts review technical SQL rules are written to patch the gap immediately. That models retain automatically, picking up on what the rules are doing to learn new patterns and apply them in a broader context.

Next step is the engine of the whole system, feature and model R&D. That's where our big detection wins originate. Once we release a new model version, we evaluate and then it's back to monitoring. A few things to call out. First, we don't only act on new reviews. When we learn a pattern today, we go back to past reviews and reevaluate them against the updated model, which happens regularly as part of our weekly retraining cycle.

Second, our research pipeline runs on two tracks. The first track is focused R&D, driven by new patterns we've observed in production. Something new shows up, we built a targeted defense for it. The second track is continuous improvement, staying current with academic research and evaluating new architectures. Many of those research projects don't stay as research, they make it into production.

For example, our research project from 2025 on coordinated fraud via sequential pattern sharing in graph. The core outcome of that research project delivered 3 times the detection volume of the prior model version. If we take a step back to look at where we came from and where we are now, it's been 12 years of evolution. We started with juristic in 2014 and built our legacy classifiers, including NLP models a couple of years later.

From 2022 onwards, you can see a rapid development in the application of more advanced technology, including the release of our first pretrained transformer-based model in January 2022, senior as the first release of ChatGPT built with the same technology. The numbers from the past four years tell the story, roughly 3.4 times the model complexity of where we were in 2022, measured across feature counts, learn parameters, free splits, graph size and roughly 3.5 times increase in filtered fraud volume over the same window, showing the effectiveness of technology in action.

A few words on the people behind this. We have 10 specialized degrees on the team, four PhDs, six masters, every one of them with years of experience in AI and machine learning. As a matter of fact, 62 years of experience total. Our team's objective is not just to keep up with the field, but to be thought leaders that actively contribute to Frontier applied research. Last year, we presented our Graph Neural classifier work at Machine learning with Europe, a leading conference dedicated to real-world AI applications. We are planning to continue with more conference contributions this year.

Which brings us to the question I get asked a lot, doesn't the AI era make all of this easier for fraudsters? And for one specific access, it does. Writing is cheap now. But text was never the strongest fraud signal anyway. What you're looking at is our actual feature space. These are broader categories that encompass some of the examples you saw in the animation earlier, colored by the signals they represent. The gray cluster labeled content, that's the entire surface area LLMs commoditized, less than 5% of our feature space.

As you saw earlier, we use AI as a tool against fraud, one of many in our toolbox. And we have a research pipeline that continuously evolves our defenses to stay ahead of frauds evolution. So while LLMs are a risk factor, they are not a game changer. We have the tools and the process to adapt just like we do with every other new fraud pattern. You've seen how our fraud detection technology operates and the outcomes it is able to achieve.

Next, Dom will cover regulation.

Domonique Rai-Varming Trustpilot Group PLC - Vice President - Trust & Legal

I'll start by saying thank you, Sona. Well, I'm here to bring the outside world in and talk to you about regulation. So I'll just start whilst hopefully the slide can catch up. When people hear the word regulation, they often roll their eyes and maybe think of constraints and cost and complexity.

But I'm here to help reframe that because for us at Trustpilot, regulation isn't friction. We see it as a competitive advantage. Because at its core, Trustpilot operates in one of the most sensitive spaces on the Internet, trust. And increasingly, what we're seeing is that trust is being chased by regulators as much as it is by consumers. And so today, I'll cover three things: firstly, how the regulatory landscape is evolving; second, how we work to actively shape that environment; and finally, how that positions us for growth.

And I'll click. So I just want to start by zooming out and just taking a look at the bigger picture. This slide is just a snapshot of the regulatory environment we're navigating at the moment. It's a pretty complex web of laws across the UK, the EU and the US. You might not be able to see it very clearly on the slide, but much of this legislation has been enacted in recent years. There's been a marked uptick in regulatory activity. And what's also changed isn't just the volume of regulation, it's the direction of travel. And I'll come on to this shortly.

What you also don't see on screen is the movement within the EU at a member state level. Many of you would have heard of the EU's flagship legislation, the Digital Services Act. This act was intended to create a unified single market for digital services. But despite those top-level EU rules, individual nations seem to be going their own way. And what we're seeing in our space is Italy and Spain have just implemented their own specific national laws to deal with reviews.

And these aren't just minor tweaks. They're raising the bar quite significantly. They're doing things like introducing mandatory ID verification for reviewers. They're also introducing really strict time limits for when reviews can be posted. We know that other countries are watching very closely with what Italy and Spain are doing, and there's a really high likelihood of a fast follow.

And so, for us and for many others, this presents quite a unique and significant challenge because a fragmented regulatory map is harder for everyone to navigate. And it reinforces to us why we can't afford to be passive here. We have to be in these rooms, and we have to be in the center of these conversations. It's our job to help policymakers understand how to protect consumers without breaking the very platforms that they rely on, and that's something my team is heavily focused on.

Okay. We've had a quick look at the global horizon, but there are some key themes coming out that I'd like to share with you. But before I get there, I also just want to draw your attention to a quieter, more subtle, but very significant shift that we've seen. With the new regulation that's come into play, we've seen a move from platforms being seen and accepted as passive hosts to platforms now being expected to be responsible actors. And so, the days of sitting on the sidelines and claiming neutrality are over.

Taking it back to this, we can distill the themes that are relevant to Trustpilot from a regulatory perspective into three key buckets: fake review detection, accountability and enforcement at scale and transparency. Taking each of those in turn. When it comes to fake review detection, it's no longer enough to react to reports. We're required to proactively find fake reviews, manipulation and business misuse, and take action before it reaches consumers.

On accountability and enforcement at scale, there's a massive drive for speed and consistency. Regulators expect that. Best efforts aren't sufficient anymore. Regulators want to see platforms taking immediate action across the board. And in terms of transparency, regulators are expecting more of a glass box approach. So we have to be able to prove with data that our processes are working, how they work and the actions that we've taken.

Now it might be easy to conclude that regulation is an impending threat to growth and innovation. And for some, it might well be, -- but for us at Trustpilot, we see this as an opportunity because for years, we spent time building the technology that Sona has presented and the governance structures to meet these standards. And the shift in regulation validates our model and our approach.

As you can see on the slide here and as my colleagues have already explained and spoken about today, our model is built around open access. We're an open platform. Active moderation, we have that in place. And we're transparent. We're open about our processes, our rules, and we showcase the actions we take with tailored trust signals.

As I mentioned before, at Trustpilot, we don't just respond to regulation. We help to shape it. We engage directly with regulators, policymakers and industry bodies across all of our key markets. What does that include? That includes work like providing input on draft regulation and legislation, sharing data and insights on review integrity and helping to define what good looks like in practice. And this really matters because the regulation, although there's lots of it already, is still being written. And our goal is really simple. We want to raise standards across the industry in a way that rewards companies who are already doing the right thing.

Another way in which we do that is through collaboration. We're a founding member of the coalition for trusted reviews. It's a space that brings together platforms, businesses and stakeholders to share information and define shared practices. And that's really powerful for us because trust is the foundation of our entire ecosystem. If consumers feel the foundation of trust is cracking, everyone loses.

But when we collectively raise the bar and standards of integrity increase across the board, consumers can search and shop with more confidence. Businesses can see higher returns and platforms like ours become essential infrastructure for global commerce. To protect this ecosystem, we've been encouraging regulators to look beyond just platforms and rules and focus on broader education because for regulation to truly move the needle, businesses and consumers alike need to be clear on their rights and responsibilities because impact only happens when everyone knows the rules of the road.

And with that, I'm going to bring this to life with a quick quiz. I don't have a slide though, but the quiz is called, is it legal? And today, when people ask questions like, can I choose who to send invites to? Can I invite my mom to leave a review? The answer to that is increasingly complex. So we're going old school which shows hands. The click will let me.

Right. We're just doing three scenarios. Scenario 1, the incentive. You're a business owner. You want more reviews. So you offer a 10% discount or a coffee voucher for anyone who leaves a review. It doesn't matter to you if it's a good or bad review. You just want feedback. On a quick show of hands for yes, is this legal?

Okay. Well, it's very suspenseful. There we go. The answer is, it depends on where you're standing. So in the UK and US, and the EU, it is legal provided you disclose that an incentive was offered. But in the US, the FTC has taken a much firmer stance and it's a hard no. And actually, in the US, the FTC have mandated that \$53,000 fine penalty can be applied per violation. At Trustpilot, our guidelines go further than some of these laws. We don't allow businesses to incentivize reviews.

All right. Scenario 2, inviting only happy customers. You're smarter than average, that's a given. You only send your Trustpilot invite link to customers who gave you a thumbs up in your private survey. If they were unhappy, you routed them to your private customer service team instead. You're just managing your reputation, right?

Quick show of hands for yes, is this legal? On the fence, okay. It is not legal. This is strictly prohibited. So this is called cherry-picking and it's illegal everywhere, and the penalties are actually much more severe than for the violation I just mentioned. So we're looking at in the UK under the DMCCA, fines of up to 10% of global annual turnover, 4% in the EU and the \$53,000 fine in the US.

Cherry-picking is an industry-wide problem. And as I said a minute ago, the consequences can be really severe here. At Trustpilot, we use a combination of technology, people and community, as you've heard, to try and find the businesses that are only inviting those satisfied customers. And when we find them, we apply things like consumer warnings like the one you saw earlier. And this is an area where we think greater regulatory education is needed, and we're very proud to be transparent in the actions that we take here.

Final scenario, calling this one the product pivot. So you sold a great pair of socks, and you managed to get 500 5-star reviews for them. But now you're launching high-end vacuum cleaner. To give it a boost, you move those 500 reviews over to the vacuum cleaner page. It's the same brand. Quick show of hands for yes, is this legal?

Solid, thank you. Let's see -- It is not. So this is called review hijacking or catalog abuse. So it's a strictly prohibited misleading practice across the board. We've got systems which are designed exactly to detect this sort of behavior if the content doesn't match the product category.

That was just a quick counter through of some of the scenarios and issues that we're seeing coming through the regulatory landscape. But what does this all mean? I can distill this down to three key takeaways. Firstly, the barriers to entry are rising. I think the days of anyone starting a review site in their garage are long gone. To operate a credible platform today, you need to have sophisticated AI moderation. We need robust governance frameworks, and you need a direct seat at the regulatory table. And I think in a high regulation world, scale isn't just an advantage, it's a requirement. And this shift naturally favors the trusted, more established players.

Secondly, trusted value. As standards rise, businesses increasingly value platforms where consumers believe in the content and external stakeholders have confidence in the processes. That strengthens our value proposition.

And thirdly, reducing risk. We're working to derisk our future. By choosing to lead the conversation with governments and agencies early in our key markets, we're moving from being regulated to being the benchmark. And this proactive engagement removes the fear of the unknown. We don't want to wait for the rules to change. We want to help write them. So in short, we're working to turn the tide of regulation into a competitive advantage for Trustpilot.

And with that, I'm delighted to lead you all on to a 20-minute break.

(break)

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

Good afternoon. I'm Hanno Damm, I'm the CFO, and I have the great pleasure to have a conversation without slides with Ciaran, our Chief Product Officer, who's been in the business for a little over a year now. It's been a fantastic impact.

And we're going to talk about the -- why openness, and we've talked a lot about why the platform is open and why that wins in the age of AI. But Ciaran, before we start with that, let's take a step back. We've seen the loops, the flywheels. Explain to us how you, as the CPO, look at these flywheels in the business.

Ciaran Dynes Trustpilot Group PLC - Chief Product Officer

Yes. Great. Thanks, Hanno. And just to set some context because the title kind of reveals all. It's just about that openness. So much of AI and large language models today are based on open content. And so, it's just important to keep that context as we think about some of the things that we're going to talk about.

In relation to the two loops, I think hopefully some people in the room have used Trustpilot in their lives and kind of have a developing opinion of when you use it. A lot of it has to do with when you're kind of in a buying journey and you kind of get to that close to maybe making the purchase of a service of a product. And you might just want to triple check your facts and go on. Trustpilot looks at what other customers, consumers or users of a particular brand have used.

And then you might find yourself either leaving a review. Now personally before I joined Trustpilot, I won't mention the name of the brand. It is a wearable device, of which there are many. But it was that kind of notion of having a terrible experience that Trustpilot was like a last port of restitution, and you can leave a review, comment on the particular brand.

And hopefully, somebody after many, many e-mails that you've tried, will see the Trustpilot review and respond back. And that is a journey that lots of people go on is that when you're just at your width end, Trustpilot is a place that you can go and kind of leave a review and maybe that particular company will respond.

But there is another journey, and that's the one where you find either small businesses or small companies, you've had a great experience. You just want to help them and you want to leave a kind of a complementary review and lots and lots of people do that. And we know that helps, right? Other consumers, other customers go to Trustpilot, they see nice things that people say about brands and indeed the bad things. But they kind of match that together and they kind of see that. And that's that first loop.

The second loop then is businesses inviting their customers to leave reviews. And what that generates for them is lots of insights. They read the reviews, they reply to the reviews. It may make them kind of think about making some changes within their business. They go through that particular loop, they measure and see whether or not that has been successful.

And then what they do is they'll tend to take some of those comments that they've got from the customers, and they will showcase them is the phrase that we use, which is putting things into their marketing materials, into the website. In some cases, I've seen

customers in their follow-up e-mails to existing customers of theirs, just as a reminder, maybe they're in a subscription and you say -- hey, just to let you know, people who have also purchased from us have said these following things. And sometimes it's just during that renewal phase, it reminds people why they should continue to subscribe.

All of that basically is the second loop. And that grows and grows and grows. It's a continuous evolution. And it feeds in that second flywheel where more reviews on the platform due to invitations, more insights drives more marketing and the thing basically catches fire from there.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

So you speak to a lot of customers and maybe walk the audience through what are you learning from them and how they use Trustpilot and what that evolution is in that customer value proposition.

Ciaran Dynes Trustpilot Group PLC - Chief Product Officer

So over the last year, kind of joined the business, I was just curious. I was deeply curious as a product person as to why somebody buys the product and why would simply buy into the proposition. I've met some big, big enterprises trying to understand what they do. And what's been interesting is the journey is so consistent and so similar.

So I've met an amazing solar panel company in Spain. They installed solar panels. It's -- Spain still has sun, lots of it. They have lots of solar panels. But in the journey, it was like, so why did you buy Trustpilot? And the story was, well, we started with the score and so many companies will say that.

And then it became what next? Well, we started to read the reviews and you go funny that. And it's so funny. They actually -- you'll consistently hear that kind of statement of we started to read the reviews. And then we thought, hey, there's something is wrong in our business, we should go fix those things, which created an okay or a goal for one of their VPs. Then that resulted in making some changes and fixes.

And then the curiosity of the executive team was, did we fix that? Are we still seeing that feedback in our reviews? And the answer was no, we had fixed that. Which then begets to kind of a reporting rhythm where their executive team, their leadership team is meeting some cases, weekly, some cases monthly, to hear what's the latest trend? What are the latest things we're seeing in the reviews.

And at the end of that little spiel that I got from this particular CEO was like, what's the impact on that in the business? He said, well, there's two things. One is every single person on my team, that was their particular team on the call, their annual bonus gets paid out based on five things. And one of those five things was Trustpilot.

So I was like, wow, it's pretty meaningful. And so when you hear that you're like, wow. It's like, well, what's it all for, because we want to be the best customer experience in class for solar panels in Spain. I was like, oh, so what do you need from me? It's like, what we've learned is the score of our brand isn't necessarily the thing that matters.

What matters is the plumbers and the suppliers and the installers of solar panels, which has got me thinking like, you got this amazing technological product, and it's down to the person who goes and installs the heat pump or the solar panel at that business. And if they screw that up, what happens is you blame the product brand, which is that CEO. He's like, that's our problem.

So I need the score of every one of my suppliers to then institutionalize such a high quality of installation that my brand does not get negatively impacted. And that story, I think, is universal. I hear that so many times. And of the ones I meet, it's that start with a score, but it evolves into this incredible rich feedback loop of continuous improvement of the business.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

Yes. And there's a couple of interesting things. We -- the company bonus component, we obviously live this at Trustpilot as well, as you well know. And we also read our own reviews. And if you read our own reviews, I think the 2 points of contention that we often see are on the consumer side, they're unhappy if we remove a review from them and they had a genuine experience, the false positives we talked about before.

And then on the business side, they're unhappy if we don't remove a review that they would have liked to have removed. So how do you, as a product person, think about that tension?

Ciaran Dynes Trustpilot Group PLC - Chief Product Officer

It's very real. Like I think the team obviously brought everybody through some of the key aspects of what we do. And it's very true to heart of the principles that we operate under like things are fair, they're neutral, open, relevant and transparent, and we live and breathe that every day. And we literally make no apologies for some of the decisions that we make on leaving a review on the platform because we believe it is a real review or we basically see no reason why it should be taken down.

And then equally, if somebody is like thinking like, hey, well, this review basically is not representative of my business, but we leave it up. And that process is the one we observe, and it is fair and it is unbiased, and we continuously refine the rules in terms of how we do it. And you heard that from the team in terms of moderation. I think one thing, again, just back to some of the customer stories is maybe it's not to dispel a myth, it's just to kind of represent what can happen.

So last week, I met with an online gardening company in the U.K. Those basically who love gardening, and I do, like Monty Don is a personal kind of -- I think I was going to say Monty Don is my favorite Britain. Is that okay to say in Britain, but he is. I love Monty Don. I think he's just an amazing -- he walks like -- back to the story. But I met this garden online company. And what was interesting is when they get a 1-star review, they walk that down into the shop floor of the business.

Now it's not there as a stick to beat people with, just to kind of mislead you, but they wanted to say is like, hey, we had this review, and it was about moldy seeds. So why were we shipping moldy seeds? Now if you are a gardener and you've seen what's been happening in the otherwise known as the British winter and spring, it's the rainy season now recalled. It has been unseasonably wet and damp. And the -- one of the team members said, we received a bunch of seeds in from one of our Dutch suppliers.

And they noticed they were subtly damped. And even though they were in the date for the seeds being still valid, there was a concern, should we ship these out? And of course, they did. A couple of days or a couple of weeks later when people received those seeds, they are moldy and those moldy seeds you cannot use. And it was just that decision-making that said, what do we want to do the next time that happens?

And it was empowering the team to go, you shouldn't have sent them. And that's so powerful. And that's what I mean by that continuous loop. You get that with Trustpilot. And similarly, just to kind of close off the good news story, when somebody's name is used as a such and such fixed my issue, they also walked down to the team and go, hey, Hanno, somebody mentioned you in a review, and they all give each other a little clap, right? So I thought it's a nice way of thinking about what happens.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

That's great. And so let's pivot a little bit to AI now and why openness becomes even more relevant in the age of AI and how these models actually digest our data.

Ciaran Dynes Trustpilot Group PLC - Chief Product Officer

So as we started, we said it was about so much about open content. There's kind of three things, we here call it the three oars, but just to kind of put it in context that so much of large language models and AI is based on relevant content. So I'm sure you've all done it. You type into OpenAI, Anthropic or Gemini, whatever your tool of choice is, a query and the content that comes back, it's relevant to that query. That's what large language models have revolutionized.

What's also interesting is the more recency of the content. Large language models are tuned to try to find new content, dynamic content, the latest content all the time. You see that as part of the way the strategy works. So you've got recency as kind of a key driver for that. And the last one then is the ranking.

So AI large language models use AI search, which is built on the 25 years of SEO ranking. They actually use this as part of the way they find information. And therefore, the ranking and domain authority of the content of the site becomes really, really crucial. So in this new AI world, Trustpilot's content is at 94 as a domain authority score. That's off the charts.

That is 0.01% of websites in the world are in the 90-plus number. So we're up there with Google and YouTube and Wikipedia and these types of websites as a place where large language models consume the information and trust that information more because it has been connected by so many more websites. That's kind of what's going on.

So therefore, the content, the citations, the mentions of the information on Trustpilot is just like large language models are hungry for that information. And therefore, when you're typing your prompt, as long as it's relevant and it's recent and it has a high domain ranking, it's more likely that's how you get citations and mentions of the content.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

And so we -- obviously, we show up in all in these answers now more and more. And all this is underpinned by this massive repository of reviews that we've accumulated over the 18 years plus. And how do you think about this as an asset? And what are you going to do with this as a product?

Ciaran Dynes Trustpilot Group PLC - Chief Product Officer

It's an incredible database. And it's -- when you listen to the team and some of the things they talked about, when I first met Sona and the team, it was kind of funny they were like you know Sam Altman. I said I don't know him personally, I wish I did. But when you think about the history of large language models and transformer models, the AI team at Trustpilot was doing the same work.

They were building our own AI language models six, seven, eight years ago. At the time, it started off with blasphemy laws and these types of things that were kind of key to do with some European legislation. But now you fast forward where it is today, every single business in the world is using transformer logic, which we know is large language models.

So you have this incredible database, but the curation of that is amazingly complex. So if I was to say the word to you something is too expensive, you might say, well, that's time to value or you might say it's priced or cost too much. There are all these kind of like synonyms of words that we kind of take for granted as humans.

And then when you factor in that's just in English and you throw in Dutch, as Thomas mentioned, you have German, I have Irish somewhat. And you mash all that together, that database is incredibly valuable. And it can be used for peer reviews, marketing research, different ways for businesses to understand what's going on. But the ability for it to zero in on what's broken in your business, it's actually harder than you think because people use different ways and different nouns and different words to describe the same thing. And we've effectively been curating that for such a long time. We have this thing that you can now embed into a business.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

Yeah. And I remember a story where we helped a company that was sort of doing marketing as a vegan -- for vegan products. And in the review context, all the reviewers were talking about meatless or meat alternative and no one used the word vegan. And so I think that's actually another real valuable insight.

Now this is about trust and not about the product, but it's in the have you. We may do a more extensive product teaching at some point. But can you give us a little teaser of how you're thinking about evolving the product in the age of AI in particular?

Ciaran Dynes Trustpilot Group PLC - Chief Product Officer

I think where we'd like to take things, we've seen some of our customers do this with Trustpilot is how do you take that database and how do you embed that into every single action every day. And those examples where we talk about somebody walking down into the building and saying, who got this one star review, can you tell me more about that? There are another group of customers that have automated that thing to an nth degree, and it's like eye wateringly good.

So I won't exactly talk about what we're going to do on the product side. I just think the opportunity that you could just imagine that if every single day, every decision you make that happens with the customers, whether that's delivering things, changing the offerings, even when you launch a product, you could come back and go, how did we go with the launching of that new product?

That was supposed to fix the solar panel installations. Did it work? When you're dealing with thousands and thousands of businesses and it's over a long period of time, how do you put that together? And I'm not going to say anymore, but that's what we're going to come back and talk about.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

And how do you think about the world of agentic commerce that's upon us at some point?

Ciaran Dynes Trustpilot Group PLC - Chief Product Officer

We are working with some of the largest marketplaces. So, I think when I look back at like 20 years ago, I remember the time that Netscape was still a company. And if you remember the time you took your credit card and you put it into an SSL, because this is all the terminology we used to back then, but you put it into a form and you click the button and you hope somebody didn't rip you off, the experience you have today was defined by Amazon's one-click patent back in 2002.

And we just take it for granted that you go on to the shop app, you click your credit cards is there and there's no questions asked. The new evolution of where we're going to go from the one click is what that's called to the buy for me, that requires a very tight loop where the vendor has some extra third-party signal that they are legitimate and that legitimacy is based on customer reviews.

It's based on a whole bunch of other factors. When you start to unpick how the Internet works today, you need to basically rewire for this new world where agents are buying things -- and therefore, where else would you go other than to ask the customer base of that brand, hey, is that brand legitimate? What's the service like? Do people get things delivered on time? Do the moldy seeds of Britain gate get restored, the big problems of today? That's what's so interesting to us, and we're going to be talking a bit more about that the next time about who are those marketplaces and how is Trustpilot embedding into those.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

Excellent. Well, that will be fun. We're now going to have Adrian, and the rest of the team come up and open the floor up for questions for some Q&A.

QUESTIONS AND ANSWERS

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

All right. Thank you. I think we're going to start in the room, and then we might have some coming in from folks online as well. But I'll just take hands up, please. Yes.

Hai Huynh UBS AG - Analyst

Thank you. This is Hai From UBS. I have a couple of questions, please. The first one is on -- so you removed 7.8 million reviews last year. But what I'm trying to gauge is where do you see the problem is? Is this kind of a ceiling or the floor of the scale of the problems that you have or you as an industry have? That's number one.

And the second one is on the cost of combating that over time as AI makes it easier to right-click reviews with scale, do you expect costs to ramp up as well as it's a race essentially for you to have better AI to combat?

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Sure. So maybe I'll take your first point and then hand to Hano for the second on costs. So 7.8 million, as we said, was a very substantial increase on the prior year. I think the way to think about that, it was over 70% more than we removed in 2024. So 7.8

million was the total for 2025.

I think the way to think about that is it's a combination. Of more aggressive, bad actors and improved detection on our part. We remove everything that our systems believe is fake, and I think we've given you a pretty detailed account today of how our systems have evolved, how our systems have got better, and how we think about the various trade-offs involved, of which there are many.

As I said right at the beginning, I think the best way to think about our effectiveness in doing this is the way people vote with their. Their behavior -- millions more people are using our platform, more and more businesses are using our platform. If you think about the examples that Ciaran and Hannah were sharing earlier, the companies wouldn't be doing this sort of root cause analysis of were the seeds moldy or what words are people using to describe vegan food or solar panels or whatever.

They wouldn't bother doing that if they didn't fundamentally at some level trust the content that is there. So I think the most useful way to think about the efficacy of all of this detection activity that we do is, are people finding the platform more useful or less useful over time? And all the evidence, as we can see, is that people are finding it more and more useful. And we are getting, every year, more and more sophisticated in how we deal with the various threats that are out there. So maybe on cost.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

Yes. So I mean, I think both Thomas and Sona illustrated very well how writing is only part of the detection. And so AI writing is actually cheap, I think is what Sona said, but everything else is still expensive. And that's where the vast majority of our fraud detection actually and our fake review detection resides.

And over the years, we'll continue to invest more and more into that. The teams are growing. But with the use of technology and all these sophisticated models that we've developed internally and are continuing to develop, we continue to see that this is not growing faster than revenue. And so we're very much unable to generate operating leverage even in that area.

And AI is helping us certainly deal with a lot of the manual parts of this in the sense that people are -- consumers are flagging reviews and someone has to take another look at this. Someone has to read this now. We can translate it, for example, out of any language into English. And so we only need English-speaking agents to actually then read it because the AI has translated it. And so there are a lot of ways that technology is actually making this more efficient for us.

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Any other questions in the room? Yes.

Timothy Ramskill Bank of America - Analyst

Thank you. It's Tim Ramskill from Bank of America. A couple of things from me, please. I guess maybe it's a bit tricky, but I thought perhaps if we use the situation of the Italian regulators kind of review into what you guys have been up to. And I thought -- because I felt that the regulatory section, it felt like you got all bases covered off the sort of the three kind of key considerations you talked about.

So when we look at that chart of all those different regulatory considerations, are the regulators understanding what they're trying to fix for? Is that -- I mean, again, I appreciate you might recap what you say, but is that part of the challenge? And again, I know you came back very robustly to the Italian regulator and said, look, we don't think there's an issue here. We're going to kind of keep pushing. So maybe just to sort of talk around that regulatory discussion, but if you wish to use that as the example.

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Sure. So I think a couple of points on that. First, we share the same fundamental goals as the regulators. As I think Dominique pointed out really well during her presentation, we have been for years calling for the things and agitating for the things that have since become laws in the various jurisdictions where we operate.

So we have exactly the same goals as the regulators in terms of consumer fairness, transparent information, all of the principles that we've outlined today are exactly the kinds of principles that you will hear the regulators propounding.

Now as you say, we are contesting rigorously what the AGCM have said and we're appealing the decision. And we believe that when the nature of how we work is really deeply understood, then that process will come to the right conclusion. But I've got nothing more to say about that at this point.

Timothy Ramskill *Bank of America - Analyst*

And then my sort of second related question is, again, I think you, in the same section referenced the kind of dynamic around sort of verification before reviews can be left, which does sound like something of an impediment to consumer engagement. So again, just interesting thought on could that likely take hold? Surely that must be a concern of some degree because it will surely slow things down.

Adrian Blair *Trustpilot Group PLC - Chief Executive Officer, Executive Director*

Yes. So that's not the way we work at the moment, to be clear. We have the functionality. It's there in the product. Some users verify themselves with that functionality, but it's not something that we are implementing across the different markets where we operate.

And I think that if you look at the trade-off involved, implementing those sorts of barriers would result in fewer people choosing to leave feedback. And we think there's huge value in operating an open customer feedback platform where we keep the barriers as low as is consistent with the sort of work we do that you've seen today to keep the platform trusted such that businesses get as much insight as possible.

And we think that's ultimately in everybody's interest to hear from as many consumers as possible. So that's the way we work across our market. And I think if you look at the key regulators in -- we talk about the European Commission, the CMA in the UK, the FTC in the US, they're very much on the same page on that stuff as well.

Hai Huynh *UBS AG - Analyst*

And I think just sort of -- because I think you picked up on something that Don had mentioned in her remarks around the pending legislation in some of the markets that are being discussed. I think as Adrian pointed out, we would not be supportive of that. And it's a question whether that's even like sort of compliant with EU Omnibus directives, et cetera. So jury is out there.

Sean Kealy *Panmure Liberum Capital Ltd - Analyst*

Afternoon everyone, Sean Kealy, Panmure. Just quite a sort of wide-ranging one really. You guys accept reviews in sort of a very broad range of languages. Just a sort of very open question. How does that complicate some of the models and approaches you use for assessing sort of what content can stay on the platform or not?

Thomas Vermeulen *Trustpilot Group PLC - Director Fraud & Investigations*

I've got this mic. Is it working? Do you hear me? Yes, because I can't hear myself from here. So this is where it comes down to the content piece, right? Like if we look at how do we detect the fake reviews, then actually the content piece is just a very small piece, and that's also what Sona tried to visualize with this around 5% of the data signals. That industry, that's content.

So -- when I look at fabricated reviews, I even don't look at the content. I look at how is this review written and how did it come on to the platform. So we look, for example, at the connectivity proxy usage. We look at what kind of device are you using. So for me, that language piece or a fake review -- fabricated review piece is actually not so important. Of course, you can get insights out of it, and we will use it, but it's absolutely not the main feature. That's why 5% of the content is only there to support the detection.

Adrian Blair *Trustpilot Group PLC - Chief Executive Officer, Executive Director*

I think it comes back to that sort of tip of the iceberg concept that we were talking about and how the vast majority of the relevant information lies beneath the surface of that iceberg.

Sean Kealy Panmure Liberum Capital Ltd - Analyst

If I can just. Just to continue on that threat. I guess the piece I'm interested is less actually about fake reviews here, more about applying it to the sort of content that you will accept. And so assume that the stuff that comes through is accurate. I'm just really interested to hear about how you go about taking the language issue.

Maj Santhakumar Trustpilot Group PLC - Senior Director Trust Operations

Yes, I can say that. We'll go through it in one of the breakout sessions to talk about how our community play a role here in reporting content as well. But Hanno touched on it earlier on. We actually use AI in our translation tools as well, which is really accurate. We validate a lot of the results as well.

We do have certain local language expertise across the teams. So it's kind of high priority for us to make sure that we're making accurate decisions alongside our policies that exist across that match to our guidelines for our reviews.

Ross Broadfoot Rbc Capital Markets - Analyst

Hi, Ross Broadfoot from RBC. The content piece, as you've said, is small, but it is big for the perception of the platform and certainly for those of us looking on that. Do you or should you have any standards for the text of a review? I can have had a 5-star experience, but not been feeling very articulate that day. My submission may look fake to other people and undermine the credibility of the platform. How is that a line that you tread?

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Yes. So we have to tread a fine line here because we don't want to put ourselves in the position of being the kind of editorial standards body for consumers as to what constitutes strong content. But the sort of thing that we do is we -- as we discussed today, we detect what our systems regard as being harmful, illegal content to get the worst stuff off the platform. We have guidelines for reviewers that encourage people to write longer, more informative content. We have a flagging system, both positive and negative. So you can thumbs-up something you find particularly useful or you can flag something that you find concerning that you think shouldn't be on the platform.

But in general, we take an approach whereby we're not putting ourselves in this kind of editorial position. And to be honest, I think that's one of the things that consumers really like and appreciate about the platform that it has that incredible kind of authenticity to it. And what we don't have a Trustpilot, which I think is an important distinction with social media, is the concept to follow us.

So if my mom and Elon each write a review on Trustpilot, the system treats them in exactly the same way. And I think that is a really important distinction between the way we work and the way pretty much all of social media works in giving everybody a voice.

Ross Broadfoot Rbc Capital Markets - Analyst

Thanks.

Unidentified Participant

Thank you very much. Thanks, Adrian and all the speakers. I've got two. And just to continue with this trend. Of the 7.8 million reviews that we've kind of removed, how many of them were, dare I call it, real or genuine? In other words, what is the accuracy of this machine?

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Yeah. So we think about this trade-off constantly between what we call false positives, which we referred to in the presentation and the opposite, which is this false negatives. And we're constantly trying to widen the space between those things. We do not publish or publicize what ratios we're working towards at any given moment because for obvious reasons, I think that would give a strong signal to people trying to misuse the platform. And it's also -- it would suggest that we think anything is acceptable, which we do not.

So when we detect, we remove and we're very -- as the team have outlined today, we're very sort of decisive in our approach there. But the effect of what we're doing with technology and the advances mean, as Sona said in her presentation, I thought very well, we are going well, well beyond the sort of academic benchmarks that you might see out there.

And I come back to the point that the best way to actually judge as an outsider how effectively we're doing this is how useful people are finding the content over time and how authoritative it's regarded by the likes of LLMs.

Unidentified Participant

I've got a second one, but I'll catch Patrick offline.

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Yeah, we will have breakouts later, of course.

Dan Ridsdale Edison Investment Research Ltd - Analyst

Thank you. Dan Ridsdale from Edison. On the subject of regulation, if the regulatory environment is potentially getting more fragmented and reviews are getting more regulated, I guess first question is how much you already customize your platform for individual geographies?

And secondly, if that fragmentation starts to get more significant, how well placed you are to continue with that adaptation, for example, if the US does something different from Europe to Spain, Italy and so on?

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Sure. Do you want to have a stab at that?

Domonique Rai-Varming Trustpilot Group PLC - Vice President - Trust & Legal

Yeah, absolutely. In years gone by, we did have a one-size-fits-all approach, and it works. But as you say, there is fragmentation afoot. So if Spain is looking or has enacted, Italy is looking to enact. For us, it comes down to a cost-benefit trade-off, as Adrian said before, we're moving in the same -- we want the same things with the regulators.

So if there is something in an upcoming piece of legislation in Italy that has the potential to have a positive impact across, we take everything on a case-by-case basis and see if it works with the product and we take a bespoke approach.

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

And generally, I would say we look for our product and the way we work to run well ahead of regulation. As we said earlier, we've got aspects of how we work that we have been advocates of long before it actually got encoded in regulation. And I think a really important point to bring out is that the conversations we have with businesses are often along the lines of us educating them about the regulatory framework that they operate in.

Back in the day, five years ago, it would have been us saying, look, here are the Trustpilot principles that we expect you to live by. Now we're not just saying here are the Trustpilot principles, we're also saying these are actually the laws of the land and you need to

live by them, which is very helpful to us in those conversations with businesses and generally in keeping the platform trusted. We've upped a quota.

Dan Ridsdale *Edison Investment Research Ltd - Analyst*

Thanks.

Unidentified Participant

Thanks guys for the the question. One for Ciaran, and I might be jumping the gun a little bit, so apologies if I am. You touched on agentic commerce. I think one of the ways to think about how that might evolve is to look at the parallels with when e-commerce started.

And as we know today, e-commerce is varying levels of popularity in different verticals, et cetera. I'm just interested, is there anything you can share as to where you think agentic commerce might go and where -- yes, I'll leave it there actually.

Ciaran Dynes *Trustpilot Group PLC - Chief Product Officer*

Yeah. Well, I think if you look at amazon.com today, you can find the buy button in certain geographies already. I think what's interesting is in that evolution like the Netscape evolution to the buy -- the one click, the Ts and Cs were critical, right? So if you look at Amazon Prime, I'm sure many people in the room use it, the ability to return something after 30 days and not cost you a penny was like that was revolutionary.

Because that notion of I can buy something online or somebody in my family can buy something online, and I have the ability to return the goods. One of the things that's problematic with agentic e-commerce is that that's just not going to cut it. So if you imagine some luxury cruise that you want to go on and you want to find the best price, let's imagine it's GBP20,000. It's a pretty expensive thing. You want to bring the family.

What you want to do is find the best discount online, and there's probably one or two cruise liners that you want to be that one, not the other one. And you could imagine agentic e-commerce pretty straightforward solving that particular challenge. Now the question is, did it buy for you and after the fact to come back and say, I bought the luxury cruise. And it could do.

But the question you have is if it was like an airline today, take any one of the airlines that you know and love, you wouldn't get your money back. There would be no refund. And without unpicking those kind of today blockers, agentic would fall apart, right? Because then you'd say, I don't trust it anymore because it hasn't got the right for refund. Those things have to be changed. And my expectations just in the conversations we have with Google and Shopify and some of these larger marketplaces is they have to solve those problems. So where does it start?

It starts with the merchant sign-up process. So when the merchants sign up to a large marketplace, they're going to have to adopt those Ts and Cs, right to refund, all these types of things. And they're going to have to build their, we would call it reputation, but they're going to have to build in a signal that says, my business is legitimate and it has a high quality of service. And that could be down to working with Stripe or some of the payment providers because they see the credit card transactions, which is really interesting. And also customer sentiment and feedback about the service.

And for things like luxury liners like \$20,000 payments, that's where it really has to work. Buying a coffee is trivial, but spending \$5,000, \$10,000, \$20,000 in a liner, a pension fund, an investment fund, you start to get to very interesting territory, and those are complicated purchases.

Therefore, you're going to have to unpick the Ts and Cs for this whole thing to evolve. So that's where I think it has to go. I think agentic e-commerce online purchasing, the Temus, the Sheins, all those brands, I think they will solve it pretty quickly. But the -- is the customer satisfied and happy piece is one thing we think is really interesting to that loop.

Unidentified Participant

Brilliant, thank you.

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

At the back there.

Thomas Brown Premier Miton - Analyst

Thanks, it's Thomas Brown from Premier Miton. Does Trustpilot have regulatory risk in the event of bad action on the part of your customers?

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Can you elaborate?

Thomas Brown Premier Miton - Analyst

So supposing one of your customers is cherry-picking, is that then liable for a fine or you for it happening on your platform? Thanks.

Domonique Rai-Varming Trustpilot Group PLC - Vice President - Trust & Legal

Well, it's a good question. And as I said, it kind of depends where you are. But yes, it is a shared responsibility in many ways. And so for us, in the instance of cherry picking, as I mentioned, it's an off-platform behavior. So from our perspective, our focus is on ensuring that we're really clearing business expectations.

We're communicating clearly through our terms and guidelines. And when we're selling our product, we're setting out those expectations too. And we have the systems in place to detect that and feedback. So, yes, but for cherry-picking specifically, the fine as drafted in the legislation at the moment applies to the business, but there will be -- there is an expectation on platforms to be doing work there as well.

Thomas Brown Premier Miton - Analyst

Thank you. And then with respect to your comments about it's not enough to be passive anymore. I imagine there are some other review aggregators who are pretty passive at the moment. Do you -- as you see the wave of regulation, do you see them having to shut down if they can't be bothered to invest in staff because it's a rather modest part of a large business?

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

I mean, look, I think the overall way we see it is that we're on the same side as regulators. And we've always, as you can see, taken extremely proactive stance to the stuff. And we've -- year-by-year, we are becoming ever more sophisticated at how we set these patents. The fact that we have so much data is what enables us to be very, very good at that.

And I would say, yes, it will be extremely challenging and wrong, frankly, to be a passive actor in the face of all of that because ultimately, this is about people and giving them tools, the useful information they can trust. And I think it will be wrong to expose consumers in the way that it would if platforms were completely passive in the face of all that. Obviously, it's not for me to comment on other businesses, but this is why we take such an active approach because we think it's the right thing to do.

Do we have any questions online?

Domonique Rai-Varming Trustpilot Group PLC - Vice President - Trust & Legal

Yes. So Jeff Pocket, Peel Hunt has asked, what happens if someone disputes a review taken off the site? What is the process for that? Second, given your product ambitions, is your second product team rightsized for the road map? And thirdly, how do you solve the cherry-picking issue of review invites? Or would it be more about better education so companies self-regulate more?

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

So maybe on the first and third points, I hand to Maj and then Hanno for the second, if that's all right?

Maj Santhakumar Trustpilot Group PLC - Senior Director Trust Operations

I can take the user appeals piece. When a customer or a user of our platform disputes a decision, that will come through to our content integrity team. We've got a team of internal experts who will assess each case independently against all of our guidelines. We use technology to assist in the decision-making depending on the reasons that it's flagged for. But ultimately, it will be made by an internal expert making an independent verification on the outcome.

I'll pass on to you for that.

Hanno Damm Trustpilot Group PLC - Chief Financial Officer, Executive Director

So when it comes to the whole invitation approach, that's where there are multiple pillars that are very important. It's about educating the businesses, and that's what we do through our policies and guidelines. There is more that we believe that we can do as an industry in total. We, as Trustpilot and the Coalition for Trusted Reviews, we believe that we want to educate more and want to get authority involved, but also social media platforms, payment providers.

Because we do believe that if we stand together in that piece, we can much better protect the consumers against any type of issues, whether it's invitation issues, whether it's fake reviews. And yes, there was also a part where there is a responsibility for the business, and that's why I personally am very happy for the regulatory input because it does make now very clear that in some countries, it is illegal. So that clarity is not only given now by us, but also by the authorities.

Maj Santhakumar Trustpilot Group PLC - Senior Director Trust Operations

And for the investment question, I think if you look at our capital allocation framework, the first priority is always to reinvest into the business in the right amount. And fortunately, we are blessed with the business with very high gross margin. So the incremental revenue that we're generating every year by the sort of mid-teens plus revenue growth rates that we're delivering consistently drops through at a higher rate, and that gives us quite a bit of room to continue to invest into the business.

Now on top of that, I think we're all extremely excited with the efficiency gains and productivity gains we're seeing from AI tools such as Claude, Anthropic, et cetera, that the engineering and product teams are using. And so I think not only me, but also Dave and Kieran say that we're continuing to invest into the business in an appropriate fashion to be able to deliver that ambitious product innovation agenda.

Domonique Rai-Varming Trustpilot Group PLC - Vice President - Trust & Legal

And I've got one from George O'Connor who is asking whether you can learn from how AI makes recommendations. So for example, when booking a holiday and use that understanding to figure out how AI could produce reviews that people trust.

Maj Santhakumar Trustpilot Group PLC - Senior Director Trust Operations

I can pick that one. So one, our TrustedTech team and my thoughts on what we do is we have like research. And that means that we -- every month, depending on what kind of timeline we put to it, 4 weeks, 6 weeks, sometimes 8 weeks, we do research. And those researchers are exactly the ones that we use to then see, okay, how is the latest technology being used and what kind of

information can we get out of that.

And that means also looking into how our recommendations done and how can we use those kind of things in our benefit. That have also led -- if you looked at the timeline and the growth of our technology and how effective we are, it is because we are using the technology and getting the learnings out from it, whether it's from new technology, whether it's from the things that we learned through our own analysis.

Adrian Blair Trustpilot Group PLC - Chief Executive Officer, Executive Director

Okay. I think that's it for questions. So just to say a couple of things before we go into the breakouts. So look, I opened the afternoon by talking about how trust is the most important thing in business. With the rise in AI -- if you can go to the slide, please, sorry.

With the rise of AI, deciding which businesses to trust becomes challenging, and we've outlined today how we have a critical role to play. Our platform is becoming the critical trust signal for the age of AI. Just to summarize what you've seen over the last few hours, you've heard Shazadi and the team describe a governance model where commercial relationships are independent from content decisions, not as a policy decision, but as an architectural fact built into the way the business operates.

You listened to Maj, Thomas and Sona take you inside the detection system that combines this graph network of millions of nodes, weekly model retraining, human experts who are actively infiltrating fraudsters operations and removing millions of fraudulent reviews every year. And of course, you've had a chance to see for yourself how hard it can be sometimes to determine at face value, whether a review is genuine and how we actually identify it through device fingerprinting, IP clustering and account behavior.

And then we heard Domonique talking about how regulation can be a competitive advantage and our role in shaping it. Hanno and Ciaran may clear the case for openness, talking about how many of our reviews come from people who were not invited because it's an open platform that businesses don't ultimately control. And because of that, no closed system can really produce the insights that this open system provides.

Every product we build is derived from those two feedback loops, staying honest. And all I want to leave you with today is the simple idea that all of this compounds, it gets bigger, it gets better as we get bigger. More reviews make the signal stronger. Every review process is a data point that sharpens detection, improves our models and makes the next fraud attempts harder to land. Being open makes the signal more valuable.

The choices we've made to let anyone publish a review, to publish our methodology, uphold 1-star reviews even when they're challenged by our largest paying customers is precisely what makes the platform trusted and the data something everyone can rely on.

Consumers, regulators, AI systems and businesses. The transparency is the product, and that trust signal is becoming a layer that commerce runs on. Infrastructure embedded in the tools and platforms where decisions are already being made, whether that's Google search results, AI answer engines, retailer product pages or comparison sites. Every one of those surfaces is a distribution point to share the genuine experiences that are recorded on Trustpilot.

Okay. And with that, we have three different breakouts. So those are going to take place in the room right behind my hand called Cosmos. There's another one on the far side of the building right over there. We'll help you get there, a room called Milky Way. And the third breakout is going to be in here.

Each of you has been assigned by name to one of these three breakouts. So business verification in this room, a case study about review sellers over in Milky Way and community reporting that's going to take place in Cosmos, the big room right behind me. So each of you has been assigned to one of those by name. And of course, you will all get to attend all three. So no one is missing out on anything, and there'll be a chance for further Q&A in the rooms as you go around.

All right. With that, thank you to my colleagues. Thanks all of you for coming along. Cheers.

DISCLAIMER

THE LONDON STOCK EXCHANGE GROUP AND ITS AFFILIATES (COLLECTIVELY, "LSEG") RESERVES THE RIGHT TO MAKE CHANGES TO DOCUMENTS, CONTENT, OR OTHER INFORMATION ON THIS WEB SITE WITHOUT OBLIGATION TO NOTIFY ANY PERSON OF SUCH CHANGES. NO CONTENT MAY BE MODIFIED, REVERSE ENGINEERED, REPRODUCED, OR DISTRIBUTED IN ANY FORM BY ANY MEANS, OR STORED IN A DATABASE OR RETRIEVAL SYSTEM, WITHOUT THE PRIOR WRITTEN PERMISSION OF LSEG. THE CONTENT SHALL NOT BE USED FOR ANY UNLAWFUL OR UNAUTHORIZED PURPOSES. LSEG DOES NOT GUARANTEE THE ACCURACY, COMPLETENESS, TIMELINESS, OR AVAILABILITY OF THE CONTENT. LSEG IS NOT RESPONSIBLE FOR ANY ERRORS OR OMISSIONS, REGARDLESS OF THE CAUSE, FOR THE RESULTS OBTAINED FROM THE USE OF THE CONTENT. IN NO EVENT SHALL LSEG BE LIABLE TO ANY PARTY FOR ANY DIRECT, INDIRECT, INCIDENTAL, EXEMPLARY, COMPENSATORY, PUNITIVE, SPECIAL, OR CONSEQUENTIAL DAMAGES, COSTS, EXPENSES, LEGAL FEES, OR LOSSES (INCLUDING, WITHOUT LIMITATION, LOST INCOME OR LOST PROFITS AND OPPORTUNITY COSTS OR LOSSES CAUSED BY NEGLIGENCE) IN CONNECTION WITH ANY USE OF THE CONTENT EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Copyright ©2026 LSEG. All Rights Reserved.