

When the Truth Becomes Fiction, Part II: The Selfies That Did Not Match

Proving you are a real person used to be free and assumed. It became a task, then a product, and the people carrying the most of it have the least room to refuse.

Jennifer McKinney · August 16, 2026 · 7 min read

Pa Edrissa Manjang delivered food in Oxfordshire. To keep working he had to photograph his own face and send it in, on demand, so a system could compare the picture against the one on file. The system kept deciding the two faces were not the same face. In 2021 the account was shut off for good.

He is Black. The software checking him was built on a commercial facial recognition service.

Manjang filed claims in an employment tribunal in October 2021, alleging indirect race discrimination, harassment and victimisation. His former employer applied to have the claims struck out. The tribunal refused in May 2022 ([preliminary judgment](#)) and confirmed that refusal on reconsideration in September 2023 ([reconsideration judgment](#)). Britain's Equality and Human Rights Commission funded the case, alongside a couriers' union. In March 2024 the company settled ([EHRC statement, 26 March 2024](#)).

The terms are confidential. No tribunal ever ruled on the merits. What is on the public record is a statutory equality regulator deciding the case was worth funding, a tribunal twice declining to throw it out, and a settlement.

A man who had done nothing wrong spent two and a half years, and needed a national regulator, to establish something that had never before required establishing. That he was who he said he was.

How the burden of proof reversed

For most of the history of dealing with strangers, authenticity was the default assumption and fakery was the thing that required proof. A forged document had to be exposed. An impostor had to be caught. The burden sat with whoever alleged the fake.

That has reversed, quietly, without anybody voting on it. Now you arrive carrying the burden. You are asked to demonstrate you are a person, repeatedly, to systems that will not explain themselves, and the work of demonstrating it is unpaid and not optional.

Nobody was asked whether they wanted the default flipped. It flipped, and then it was monetised.

A password is a secret you chose. If it leaks, you choose a different one and the old one becomes worthless. That property is the entire basis of password security. It fails safely.

Credential	Revocable	What that means in practice
A password	Yes	You chose it, you can change it, and a leaked one becomes worthless the moment you replace it.
A face	No	You did not choose it, you cannot change it, and a leaked one stays valid for as long as you have the face.

A face does not have that property. Neither does a fingerprint or an iris. When biometric data leaks from a vendor's system, and it has, there is no reissue. The credential is your body and it is the same body tomorrow. Failing safely is what makes a password worth using, which is why a breach at a verification vendor is permanent in a way a password breach is not.

When neutral systems are not neutral

The argument for automated verification is that it is neutral. It applies the same test to everyone. Whether it performs the same for everyone is a different question, and it has a measurable answer.

In July 2023 the journal *Patterns*, published by Cell Press, ran a peer-reviewed study of AI text detectors by Liang, Yuksekgonul, Mao, Wu and Zou. The detectors were run against essays by native English speakers and essays by non-native speakers ([Liang et al., Patterns 4\(7\)](#)).

Detector result	Share of human-written essays affected
Native English writers falsely flagged	0

Non-native writers falsely flagged, average across detectors	61.3
Non-native writers flagged by at least one detector	97.8

Every essay tested was written by a human. Roughly one in five TOEFL essays was misflagged unanimously by all seven detectors examined.

The detectors were not measuring whether a machine wrote the text. They were reacting to a narrower vocabulary and simpler sentence construction, which is what writing in a second language often looks like.

The consequence lands on students who already had the least margin. Several universities have since turned their detectors off.

Orion Newby, a first-year student at Adelphi University who is autistic, was flagged by a detector for an essay he wrote. He ran the same text through two other detectors, which reported human authorship. The university ordered him into a plagiarism workshop and told him he could not appeal. In February 2026 a court found the finding against him was without merit ([Inside Higher Ed, February 2026](#)).

What the procurement record shows instead of market estimates

Market research on the identity verification sector is produced by firms that sell reports to the companies in it. Those figures are not usable here. What is usable is the public financial record: what specific agencies actually contracted to pay, filed in systems where the numbers carry a contract number.

Public record	Figure	What it covers
Treasury blanket purchase agreement	\$1.03B	Ceiling value for identity and access management including facial recognition. Award 2032L226A00006, ordering period running to December 2030, filed on USAspending.gov.

Congressional oversight finding	\$45M	Received across at least 25 state agencies for identity verification on pandemic unemployment programmes. Released November 2022.
GAO report GAO-25-107273	150M+	Times users accessed applications requiring the highest identity assurance tier between 2021 and 2024, with the vendor as sole provider to the tax authority. Published 11 June 2025.

These are procurement records and oversight findings, not market estimates. The GAO report separately found the agency had set no measurable performance goals for the programme and had not documented the vendor's use of AI in its required inventory.

The same oversight review found that the vendor had directed between 10 and 15 percent of benefit applicants to prove their identity by video chat with staff, because the automated face match had not verified them. One in seven or eight people, waiting in a queue to be looked at by a human, because the machine could not tell they were themselves.

That review also compared the company's public claim about fraudulent unemployment claims against the Department of Labor Inspector General's assessment ([House Oversight finding](#)).

Estimate source	Pandemic unemployment fraud (USD bn)
Stated by the verification vendor's chief executive	400
Assessed by the Labor Department Inspector General	45.7

Both figures describe the same problem in the same period. The larger one was stated in June 2021 by the head of a company selling the remedy.

The estimate of the problem was produced by the party selling the solution. It ran roughly an order of magnitude above the government's own assessment. It circulated anyway.

Who is writing the standards for authentic content

There is an international body working on standards for multimedia authenticity. Watermarking, provenance, deepfake detection. It is convened by the International Telecommunication Union together with the International Electrotechnical Commission and the International Organization for Standardization, which is to say it sits inside the UN system ([AI and Multimedia Authenticity Standards Collaboration](#)).

Its chair is a senior director of artificial intelligence and data at Shutterstock. Participants include a content provenance coalition and an authenticity initiative founded by Adobe.

None of this is hidden and none of it is improper. Standards bodies include industry because industry has the engineers and the deployment reach. Excluding them produces standards nobody implements.

The definition of what counts as authentic content is being drafted by a collaboration led by a company that licenses content, alongside firms that build the tools generating the synthetic material.

The doubt arrived first. Then the remedy arrived, priced, from adjacent rooms.

Who carries the cost of being wrong

Verification burden does not fall evenly. It falls on people applying for unemployment insurance, on benefits claimants, on gig workers whose accounts can be switched off, on immigrants, on students. On people whose relationship to an institution is one where refusing the scan means losing the money, the job, or the degree.

A person with a lawyer and a salary encounters this as friction. Manjang encountered it as the end of his income, and it took a national equality regulator to get it looked at.

The scans work correctly for most people most of the time. That is why the failures stay quiet. A system right ninety-something percent of the time produces a small population who cannot get through and cannot find anyone to appeal to, and that population is not randomly drawn.

Sources

- [Equality and Human Rights Commission, statement of 26 March 2024](#)
- [Manjang v. Uber Eats UK Ltd and others, preliminary judgment, May 2022](#)

- [Manjang v. Uber Eats UK Ltd and others, reconsideration judgment, 18 September 2023](#)
- [Liang, Yuksekgonul, Mao, Wu and Zou, GPT detectors are biased against non-native English writers, Patterns 4\(7\), 14 July 2023](#)
- [Adelphi AI plagiarism decision, reported February 2026](#)
- [House Oversight, evidence on facial recognition company ID.me](#)
- [AI and Multimedia Authenticity Standards Collaboration, ITU](#)