MARLUC

CASE STUDY · MARLUC SYNOPSIS

Udbetaling Danmark fraud-control algorithms

IMPLEMENTED 2012–2024 · GOVERNMENT AND PUBLIC SERVICES · DENMARK · BAND C · PARTLY PROVED

VERSION 1.0 · PUBLISHED 15 NOVEMBER 2026 · UPDATED 15 NOVEMBER 2026 · FACTS CHECKED 15 NOVEMBER 2026

Denmark's benefits payment agency and its delivery company use up to 60 machine-learning models to flag possible benefit fraud for caseworkers. Amnesty International examined four and reported risks of discrimination and mass surveillance. The agency disputes parts of the findings.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

AI used: Machine-learning models that flag benefit recipients for fraud investigation, including a supervised anomaly-detection model.

Implemented: 2012–2024. Grade: band C, Process change, Partly proved.

Headline result, as reported: Amnesty reports that almost 90% of cases opened from one model, Model Abroad, ended in no further action in 2023; UDK disputes the social scoring finding

1 · The case

The facts, as reported by the 3 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Udbetaling Danmark (UDK) and ATP, Denmark. Danish public body that pays out social benefits. Set up in 2012; its staff and IT are supplied by the pension company ATP under an administrative agreement, per Amnesty. Canvas block: 8. Governance and risk. Also 4. Operations and process, 6. Data and technology.

Business challenge

The challenge. UDK centralised benefit payments that municipalities had handled. By Amnesty's account, it and ATP built data-driven controls to find benefit fraud and errors, such as people who have moved abroad, or who claim single-person benefits while living with a partner.

How it was addressed. UDK and ATP's Joint Data Unit links personal data from public registers and runs up to 60 models that produce lists of people at statistically higher risk of receiving benefits wrongly. Control-team caseworkers review the lists and decide which become investigations. Amnesty says ATP developed the models with private companies including NNIT. Amnesty examined four of the models, including Really Single and Model Abroad, through freedom-of-information requests and interviews. [1] [2] [3]

Who implemented it

In-house (In-house): Administers benefits for UDK and designs the fraud-control models, per Amnesty.
NNIT (Software vendor): Subcontracted to develop some of the fraud-control algorithms, per Amnesty.
Amnesty says other private companies were also involved; they are not named in the sources opened.

Technology (SEAM pillar E)

SEAM types: T8 Predictive / ML models. Deployment: Not reported. Vendors and products: NNIT, ATP, Really Single, Model Abroad.

Where it was built in: Models run on personal data merged from public registers by the Joint Data Unit; flagged lists go to control-team caseworkers at UDK and municipalities.

What changed

Before: fraud and error leads came from reports by neighbours and public authorities. After: by a UDK official's account to Amnesty in January 2024, about 1,800 of 5,000 to 6,000 control cases a year came from the models, and the team wanted more. [1] [2] [3]

What it cost

Not reported.

What it saved

Not reported in money terms. UDK's 2023 figures, as reported by Amnesty: Really Single, 135 of 412 requested cases resulted in revenue; Model Abroad, 36 of 511. [1] [2] [3]

Governance and control

UDK says every action depends on a caseworker reviewing the case and that its controls do not amount to social scoring under the EU AI Act. Amnesty says UDK and ATP did not provide detailed evidence or independent assessments, refused a collaborative audit, and that UDK's board does not oversee ATP's daily algorithmic practices. The Danish Data Protection Authority told Amnesty it acts mainly on complaints.

S2 · Accountable owner · Claimed
UDK's board is responsible for ensuring ATP fulfils its tasks, but Amnesty says it has no supervisory oversight of ATP's daily data and algorithmic practices. [2]

A1 · Safe use · Not reported
Not reported.

A2 · Human in the loop · Evidenced
UDK says any action depends on a caseworker reviewing the case, and Amnesty describes the identification and investigation process as subject to caseworker review. [2]

A3 · Monitoring and assurance · Claimed
UDK assesses quality by the share of flagged cases the control team accepts for manual handling, and says model training cases are continuously monitored. [2]

A4 · Vendor AI · Not reported
Amnesty says NNIT cited confidentiality and gave no information on its contract or human rights due diligence. [1]

A5 · Agent controls · Not reported
No agents described.

A6 · Standards and regulation · Claimed
UDK says its controls do not constitute social scoring under Article 5 of the EU AI Act and comply with EU and national law; no regulator has ruled. [2]

What worked, what did not, what is unproven

WORKED

– UDK says controls have a defined purpose, are proportionate and aim at the correct payment of benefits.

– Caseworkers review flagged cases before any action, by UDK's account and Amnesty's description.

– In 2023, 135 of 412 cases requested from the Really Single model resulted in revenue, per UDK's figures reported by Amnesty.

DID NOT

– Amnesty reports that almost 90% of cases opened from Model Abroad in 2023 resulted in no further action, and 54% for Really Single.

– Amnesty argues that the models' use of foreign affiliation and unusual living arrangements risks discriminating on national origin and migration status, and that citizenship can act as a proxy for race or ethnicity.

– UDK and ATP rejected Amnesty's requests for a collaborative audit and gave redacted documents only.

UNPROVEN

– Whether the system amounts to prohibited social scoring under the EU AI Act: Amnesty says it does; UDK says it does not. No regulator or court has decided.

– Amounts recovered and cost of the programme: Not reported in the sources opened.

– Changes since November 2024: Not reported.

Timeline

2012 Danish government establishes UDK to centralise payment of welfare benefits; ATP provides its staff and IT [2]

2019 UDK reported to be using about 60 AI and machine-learning models to identify likely benefit

2023 UDK provides Amnesty with performance statistics for its data-driven controls [2]

Nov 2024 Amnesty International publishes Coded Injustice; UDK and ATP reject a collaborative audit and deny that the system is social scoring [1]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: Not applicable. Grade: C, Process change. Proof: Partly proved. Level of change: Process change.

Process change (Level 2, band C). The models feed the selection of fraud-control cases and supply a share of the control team's caseload. No regulator or court finding exists, so the marker is Partly proved.

The SEAM view and lessons learnt

The risk sat in the design inputs and in who could look inside the system. Models that use proxies for nationality or living patterns were tested by the agency against its own definitions, not by an outside party, and the agency declined an outside audit. For any large organisation using models to select people for checks, the transferable point is to record false-positive rates by group and allow independent review.

1. Test models for proxy variables such as nationality before use.
2. Track and publish the share of flagged cases that end in no action.
3. Allow an independent audit of models that select people for investigation.

Evidence

Score 2 of 4. Independent report with the agency's responses recorded: Amnesty's analysis of FOI documents, interviews and UDK statistics; the agency disputes parts of it.

Evaluation method. Civil society investigation: freedom-of-information documents on four models, interviews with UDK and ATP staff, UDK performance statistics and the agency's written responses. Not an audit; the agency refused access to code and data.

Sources disagree. Amnesty says the system is likely to be social scoring under the EU AI Act and should be paused. UDK rejects this, saying its controls are proportionate and comply with EU and national law. Amnesty says UDK gave no detailed evidence for its position.

Status. Reported November 2024. Current status of the models: Not reported.

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/udbetaling-danmark-fraud-control-algorithms

[1] Amnesty International: Denmark: AI-powered welfare system fuels mass surveillance and risks discriminating against marginalized groups (12 November 2024)
https://www.amnesty.org/en/latest/news/2024/11/denmark-ai-powered-welfare-system-fuels-mass-surveillance-and-risks-discriminating-against-marginalized-groups-report/

[2] Amnesty International: Coded Injustice: Surveillance and Discrimination in Denmark's Automated Welfare State (report, November 2024)
https://www.amnesty.org/en/wp-content/uploads/2024/11/EUR1887092024ENGLISH.pdf

[3] Fortune (AFP): Denmark's renowned safety net turns into a political battleground as AI and algorithms target welfare recipients (13 November 2024)
https://fortune.com/europe/2024/11/13/denmark-renowned-safety-net-turns-into-a-political-battleground-ai-algorithms-target-welfare-recipients/

Udbetaling Danmark (UDK) and ATP, Denmark website: https://www.udbetaling.dk

Version 1.0 · Published 15 November 2026 · Updated 15 November 2026 · Facts checked against sources 15 November 2026 · Downloaded 26 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/udbetaling-danmark-fraud-control-algorithms

VERSION HISTORY

v1.0 15 November 2026 First published