

Smart Check — Amsterdam Built a 'Fair' Welfare-Fraud AI, Then Shelved It When Bias Reappeared in Live Testing

AI in the Public Sector · Good practice · Netherlands

Quality score: 73/100

Evidence strength: Quasi-experimental

Summary

Amsterdam spent about five years and roughly €535,000 building Smart Check, a welfare-fraud screening algorithm engineered to avoid discrimination. A live 2023 pilot of about 1,600 applications was no more accurate than caseworkers and still showed bias, so the city shelved it.

Key dimensions

Dimension	Score
Evidence of impact / measured public value	3/3
Transparency, fairness & accountability	3/3
Transferability / demonstrated replication	2/3
Scalability beyond pilot	0/3
Governance, capability & sustainability	3/3

Evaluated by C-NAPSE

Data sources

[C-NAPSE web research](#) — accessed 2026-08-20

[City of Amsterdam, Department of Work, Income and Participation](#) — accessed 2026-08-20

[MIT Technology Review: Inside Amsterdam's high-stakes experiment to create fair welfare AI](#) — accessed 2026-08-20

[Lighthouse Reports: The Limits of Ethical AI](#) — accessed 2026-08-20

Cite this — Evidoria. *Smart Check — Amsterdam Built a 'Fair' Welfare-Fraud AI, Then Shelved It When Bias Reappeared in Live Testing*. Persistent ID: 25988c02-ada0-493e-bc0e-31870aac0bf9.