



**30th IOTA
General Assembly**


**Budapest
17 June 2026**

AI-Driven Risk Analysis in Action: Experiences and Lessons Learned in Italy

Antonino Virgillito
Head of Data Science Office

Italian Revenue Agency

Data science at the Italian Revenue Agency



Risk Analysis Sector

Data scientists, statisticians and tax experts

2018: hiring of two **Senior Data Scientists**

2023: hiring of 30 junior Data Scientists

2024: **Data science office** established

Provide support to other business structures for everything concerning data analysis (compliance enforcement)

Data science for risk analysis



Rule-based

Deterministic



RISK CRITERIA

AI / Machine Learning

Probabilistic/Predictive



RISK PROFILE





Why and when do we need AI in the first place?



What is the expected business value?



What benefits did we get in the end?




Case studies:

1 - Predictive modeling

2 - Fraud detection

Case 1: AI for predictive modeling



- AI systems trained on historical data that generates predictions about a given characteristic of a taxpayer's profile

Case 1: AI for predictive modeling

What to predict?

Initial focus: **audit outcome**

Never made it to production...

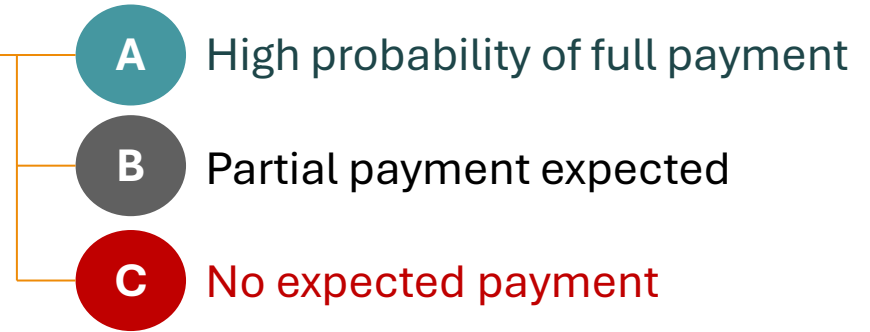
- Too many different situations not well represented in historical data
- High level of segmentation of population
- Few cases for training
- Maybe a deterministic approach is better...

Shifted focus to predicting
taxpayers' behavior

Project ISOLA: solvency score

- Forecast of the possible **behavior** that taxpayers might adopt in the event they receive a formal notification of a tax assessment

IOTA Digital workshop on AI for risk analysis
April 2026



Aimed at **supporting** territorial offices in the definition of audit plans

NOT USED for automated decisions

Results

- Predictive Accuracy (ex post): **88,7%**

- How to measure **business value**?
 - Behavioral predictions influence **process** metrics rather than **business** metrics
 - Difficult to compute measurable indicators – no right data for the job
 - Early analysis show “signals” of improvements
 - Gathered informal feedback from users

User adoption



“I trust the algorithm!”




“I want explanations!”

The problem with **explainability**

- No one *really* knows what it means
- Does not come standard with algorithms
- Cumbersome to implement

Case 2: Fraud detection

- AI algorithms for detecting high-risk companies that act as missing traders in VAT frauds

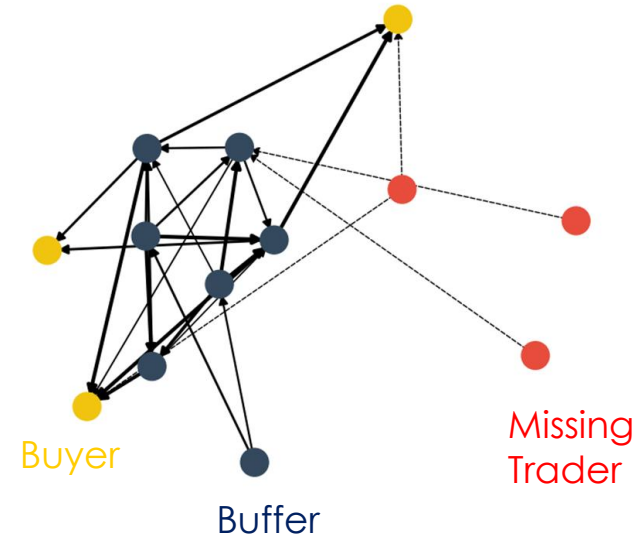


Case 2: Fraud detection

- First attempt – unsupervised statistical classification
- Feedback from users was not enthusiastic...
 - “We already detect many cases!”
- Maybe the problem is too “easy” and AI is not the solution

Proposal: Network analysis + AI

- Automated processing of e-invoice data to reconstruct transaction chains
- Detect boundaries of fraud scheme through community detection algorithms
- Supervised classification of graph nodes to predict their role in the scheme



Results

The good – it works

- Good predictive performance (around 75% accuracy)
- Output was useful in practice and well-accepted by business users

The bad – difficult to reproduce

- A LOT of manual work involved to prepare data, train/execute the model and prepare the output
- Output is rather dense and takes time to be processed by final users

Improvements



Upon business users' suggestion, the algorithm was enhanced with **text classification** of invoices



New methodology based on
local LLMs

**IOTA Forum on Tax and
Data Science**
June 2025



- Classifies fraudsters but also **tracks product movements**
- Separately detects MTs and terminal nodes, then reconstructs only most significant cases

Conclusions

- ❑ Over nearly a decade, numerous machine learning and AI projects have been developed and deployed across the Agency
- ❑ Early unsuccessful attempts were critical in helping us grasp some structural concepts and sharpen our long-term strategy
- ❑ AI should not be used for its own sake – the operational and business objectives should be clear from the start and explained to users
- ❑ Gaining acceptance of "algorithmic profiling" requires a cultural shift. Trust is built through transparency, collaboration and rigorous quality control
- ❑ Involving business domain experts right from the conception stage ensures clearly defined benefits and results in better, more practical solutions.