

# Revisiting the Risk Elicitation Puzzle: A Machine Learning Approach

Supervisor: Dr. Konstantinos Georgalos (Economics, Lancaster University)

[k.georgalos@lancaster.ac.uk](mailto:k.georgalos@lancaster.ac.uk)

**Introduction:** A foundational assumption of economic theory is that risk preference is a stable individual characteristic, a genuine trait of the decision-maker that can, in principle, be measured and used to predict behaviour across contexts. Nevertheless, evidence from the experimental economics and psychology literature shows that people's preference for risk changes considerably when measured using different methods leading to the so-called *risk elicitation puzzle*: risk preferences appear to be context dependent and different cognitive processes can lead to varying preferences. The literature adopts two approaches, a pessimistic one: risk preference is not a stable trait at all, and a constructive one: risk preference is a stable underlying trait, but each elicitation method measures it with error and method-specific bias. This raises a fundamental question: *is there a general factor of risk preference that encompasses both general and domain-specific components and can be used for the assessment of risk preference both in the lab and in the wild?*

**Objective:** Both aforementioned approaches rely either on structural econometric estimates or classical psychometric models as inputs. No study has yet asked whether latent risk preference structure is recoverable directly from raw incentivised choice data using machine learning methods. This project aims to investigate, proposing the use of both classical (PCA, factor analysis) and modern machine learning (sparse autoencoders, regularised regression) methods to discover the low-dimensional structure of choice vectors without imposing a parametric utility model. The recovered latent dimension is then given economic content by testing whether it is consistent with structural parameter estimates (utility curvature, probability weighting) from best-fitting economic model specifications, or reflects something else entirely.

**The internship:** The intern will implement the full pipeline including data harmonisation across elicitation methods, classical dimensionality reduction, and sparse autoencoder training and also conduct the correlation analysis with structural estimates (estimates already exist from past work). The internship contributes towards two deliverables: an open-source, documented software pipeline that any experimental economist can apply to their own dataset without reimplementing the methods from scratch, and an academic paper summarising the methodology and results. Through this process, the student will gain experience in machine learning, reproducible research software development, and high-performance computing workflows.

**Computational needs:** The project uses multiple experimental datasets from the experimental economics literature, comprising thousands of subjects completing incentivised lottery tasks across multiple elicitation methods. Three computational demands motivate the use of high-performance computing infrastructure: (1) sparse autoencoder training requires extensive hyperparameter search across network architectures and regularisation parameters, with each configuration evaluated via cross-validation; (2) bootstrap inference for the latent dimension recovery and its correlation with structural estimates requires the full pipeline to be rerun across thousands of resampled datasets; and (3) cross-method validation repeats the full pipeline independently for each elicitation method, multiplying the computational burden by the number of methods. The dRTP co-supervisor's expertise in Research Software Engineering will be invaluable in guiding best practices in code structure, documentation, and validation, ensuring the pipeline is robust, reproducible, and fit for open-source release.

**Tentative schedule:** Weeks 1-2: Literature review & task and data familiarisation; Weeks 3-4: Classical dimensionality reduction (PCA and factor analysis); Weeks 5-6: Sparse autoencoder training & correlation of recovered latent dimensions with structural parameter estimates (CRR curvature, probability weighting) from EU and RDU specifications fitted separately per task; Weeks 7-8: Robustness checks & write-up of results; code documentation and repository deposit.

**References:** Frey et al. (2017), *Science Advances* 3(6); Pedroni et al. (2017), *Nature Human Behaviour* (1); Zhou & Hey (2018), *Experimental Economics* (21); Friedman et al. (2022) *Games & Economic Behavior* (133); Peterson et al. (2021), *Science* (372).