

2026



Banc Ceannais na hÉireann  
Central Bank of Ireland  
Eurosystem

Insight   
RESEARCH IRELAND CENTRE FOR DATA ANALYTICS

# AI & Data Science PhD Funding Opportunities – Central Bank Frameworks – 2026 Call Document

## 2026 Call

Deadline: April 10<sup>th</sup>, 2026

## Introduction

The focus of this year's call is on supervising AI and related fundamental questions that will shape policy in the coming years. This is of particular relevance to the Bank because AI deployment in high-stakes settings, including financial services, may materially alter the risk landscape for consumers. At the same time, AI capabilities have advanced rapidly, with increasing deployment of more complex, compound AI systems, including Agentic AI (AAI) systems.

**There are two focus areas for 2026, covering the following topics:**

1. joint approaches to multiplicity, sparsity, fairness and recourse in financial services
2. developing evaluation science for complex AI systems in financial services

On focus area 1, the Bank possesses internal capability to apply existing methods to these challenges, we are seeking to develop novel methodologies through this funding and potential for collaborative research public interest. This will broaden our supervisory capacity for AI and inform policy development and implementation in Ireland in the coming years.

Focus area 2 aims to provide enhanced public interest understanding and development of approaches to evaluating more complex agentic AI systems - thereby strengthening the Central Bank's ability to identify and assess current and frontier AI risks.

We welcome expressions of interest from supervisors across a broad range of relevant disciplines, including (but not limited to) computer science, statistics, cognitive science, economics, social sciences, and related fields. We strongly encourage diversity within joint supervision teams. In particular, we welcome applications from researchers whose expertise intersects with one or more aspects of these topics, even if their work has not previously focused specifically on financial services. The Bank will provide additional context and domain-specific expertise to support the successful applicants where required.

## Important Dates

Some important dates for the 2026 call are as follows

	Call Item	Date
1	Programme Briefing to prospective supervisors	19, 20 March
2	Full Call Document available and on websites	20 March
3	Submission deadline for applications	10 April
4	Decision on successful applications	1 May

---

## Focus Areas for 2026

### 1. Joint evaluation frameworks for multiplicity, sparseness, fairness, and recourse in financial services

This topic draws on model multiplicity literature, which examines the benefits and risks arising from the existence of multiple models that achieve near-equivalent performance on a given dataset. For any particular model class, this involves identifying the set of well-performing models—known as the Rashomon set—which has become an active area of research in recent years. Several related research streams address complementary aspects of this challenge. Sparsity research focuses on identifying simpler, equally performant models from the broader solution space. Fairness research uncovers and can mitigate potential discriminatory behaviour in automated decision systems. Recourse research provides recommendations to help individuals overcome unfavourable outcomes from decision-making systems.

Recent research has begun to examine various combinations of these interconnected concerns. This emerging body of work highlights the need for joint evaluation frameworks that treat predictive performance, fairness, and interpretability as interrelated dimensions of trustworthy AI systems.

In high-risk deployment contexts such as determining access to essential financial services, these concerns all matter. Until recently, the multiple strands of the multiplicity and sparsity literature had been treated separately. Similarly, fairness and recourse have traditionally been addressed as distinct challenges. However, there is growing recognition that these issues are fundamentally interconnected and require integrated conceptual and analytical frameworks.

A research project in this area might explore several of these aspects in combination, drawing on theoretical and/or applied perspectives. Potential approaches might include development of theory, experimental strategies, causal inference methods, counterfactual explanation approaches, or other relevant methodologies from applicants' areas of expertise

For instance, one such approach may involve developing the theoretical framework applicable to various types of models, such as [Tierney et al., \(2025\)](#), or the broad set of potential paths mentioned in [Liu et al., \(2025\)](#). One [recent example by Barrainkua et al., \(2025\)](#) defines new societal level metrics to measure recourse. [Another approach characterised this as a trade-off](#) as a search for statistical optimal balance between fairness and accuracy. Other recent work on [multiplicity challenges](#) this notion of a strict fairness-accuracy trade-off and requires a critical evaluation of algorithmic choices.

A project in this area may focus on methods to find good models from the set of all models. This includes enumeration, approximations, or sampling approaches. Relating the identification of the set of good models to searching for models with specific fairness/sparsity properties is an open research area. There has been much progress for specific model types, like sparse decision trees, there has been less progress for other types. Empirical results [from Dai et al., \(2025\)](#) suggest that how one looks for fairer models can influence success in finding fair models among the set of good models. This

is important if fairer models can lead to equally good outcomes for firms deploying them and people subject to their decisions.

[Other research](#) has shown that for specific types of models, there are sometimes trade-offs for given datasets, between fairness, sparsity, and accuracy given certain types of models like scoring systems and decision diagrams. However, it is not yet clear whether this is the case for regression-based models like generalised linear or additive models or machine learning methods like sparse-decision trees, ensembles, or deep-learning methods.

This project might use the tools of causal inference applied to fairness, where there is knowledge of the causal mechanism. Work by [von Kügelen et al., \(2022\) using this causal approach](#) suggests that recourse unfairness remains a key concern on real world datasets. Another emerging approach may be to look at the impact of multiplicity on counter-factual explanations where they are actionable and useful for recourse. This could be based on any [number of approaches](#) but would have to be applicable to high risk settings including in financial services. Related to this approach, evaluating the efficacy of counterfactual explanations compared to alternatives as in [Upadhyay et al., \(2025\)](#) might be another approach to consider.

## 2. Developing evaluation science for complex AI systems in financial services

Generative and Agentic AI (AAI) represent recent developments in AI capabilities, following earlier machine learning approaches. Both have potential to accelerate beneficial uses of AI across industries, including financial services. At the same time, they may pose risks and harms to both consumers and firms. Because these systems can be one or multiple AI's themselves and can be orchestrated in several ways, this leads to a complex system with various challenges in assessing risks and potential harms.

Evaluation frameworks for Generative and Agentic AI are being developed by frontier AI labs and large technology providers industrial research groups, and in academia. These include both evaluation frameworks and benchmarks to help assess Generative/ AAI system performance. However, these are based on companies' approaches to safety and may not necessarily reflect EU perspectives regarding application safety in various high-risk contexts. This project could contribute to developing the evaluation science for complex systems with application to financial services in two main areas, as discussed below.

### 2(a) Application of statistical and data science to AI evaluation

This sub-topic focuses on applying and developing data science and statistical approaches to improve rigour in some aspects of evaluation science so that evaluators can be confident in their conclusions. Uncertainties in performance can be quantified, and improved statistical methods can be used to understand better performance, handle the specific data structures generated during evaluations, or require less test data to assess risks and harms from these systems.

Potential research directions within this sub-topic include examining the structure and approaches of current Generative/Agentic AI evaluations. For example, Generative AI

### Insight Research Ireland Centre for Data Analytics & Central Bank of Ireland

---

evaluations can be nested and hierarchical. Existing work highlights the need to take this into account to conduct valid inference. This includes adaption of existing modelling approaches ([Luettgau et al. \(2025\)](#)) to conduct valid inference, methodology that draws on design for [constructing evaluations](#), using various [information theoretic approaches](#), [modelling missing responses and the clustered nature of evaluations](#). This could consist of new applications of data science and statistical modelling [to developing evaluation frameworks](#). Agentic AI evaluations may be even more challenging to model given potential for multiple steps/sequential decision making, incomplete tasks, multi-turn evaluations, and [multiple potential outcome metrics](#) that include task completion but also reliability and harm measures.

Another approach could be extending existing work that uses Item Response Theory (IRT) models to estimate alignment/preferences in the alignment stages of model training or testing the efficacy of that training for safety purposes such as advice in regulated industries like financial services.

The stochastic nature of LLMs means that uncertainty quantification of evaluation scores is necessary (Miller, [2024](#)). For safety assessment, being able to test LLMs and agentic systems with smaller test samples may also be important (Polo et al., [2024](#)), as generating unseen test data is both human expertise and compute-resource intensive. Generative and Agentic AI evaluation is a challenge when there are sampling biases like covariate shift, selection biases in ground truth labels or judge provided labels, and missingness and clustering arising from single turn or multi-turn user conversations or agent trajectories leading to sequential sampling approaches.

Finally, when other LLMs are used as judges in evaluations, this can introduce noise to observed measures as well as certain biases. To debias these, under certain sampling assumptions evaluation datasets would amount to a substantial proportion of ground truth labels, particularly when models being evaluated are at the frontier (i.e., the judge may be less capable than the model under evaluation). Recently, progress has been made in quantifying these types of limits on scalable auto-evaluation. However, they rely on simplifying sampling assumptions. These may require more flexible regression approaches, reconsideration of sampling strategies to include stratification or other approaches, and estimation under clustered settings that arise from the evaluation design. A project in this area may want to address some of these challenges.

### 2 (b) Developing evaluation science

A second area of focus for Topic 2 involves developing methodology and approaches for GenAI or AAI benchmarks for financial services, including risks/potential harms as well as performance. Contributions to the methodology of benchmarking could include [evaluation rubrics](#), [good practice in construction and reporting](#), and how the [benchmarks are designed and implemented](#). Many of the AI-task or domain specific benchmarks are performance-related benchmarks. These are standardised, simplified measures of systems that may or may not represent real world use. Crucially, the limited existing benchmarks for AAI in financial services focus on task completion from the perspective of the AAI deployer and not the consumer or a supervisor's view of risks/potential harms. This type of work could be focused on one area considering risks and harms to consumers, firms, and wider financial stability and operational resilience concerns. A recent example of the type of approach that could be applied in financial services is by [Rabanser et al., \(2026\)](#).

### Insight Research Ireland Centre for Data Analytics & Central Bank of Ireland

---

A different approach in this area could be to combine benchmark testing with ‘human in the loop’ studies. This would be a different approach to evaluation based on combining both human and automated evaluations to assess the impact of ‘human in the loop’ (HITL) in terms of both performance and risks. This would possibly be grounded in realistic deployment scenarios within financial services. Previous work suggested that human feedback could reduce user trust in a system even if accuracy improves. However, some of these findings may be context-specific and may not generalise to financial services (like [image rating tasks](#)). Others have made detailed arguments for the benefits of [combining both approaches which would outweigh the drawbacks](#). This requires a thoughtful approach designing experimental and agent system design, including how interaction occurs, what the effects are on the human and the AI in learning to defer/escalate to a human. It would also require overcoming some of the [shortcomings of novel approaches in this area](#). In high risk or safety critical areas, it is likely that a holistic approach to evaluation would need to consider the full deployment context in which a human is in the loop.

## Application and Evaluation

The application and evaluation process consists of multiple stages:

- Call for Expressions of Interest (EOI): This document serves as the official Call for Expressions of Interest from potential PhD supervisors. The Call will be posted on the CBI, Insight, and other relevant websites.
- Submission Guidelines: Detailed submission instructions, including deadlines and the submission portal, will be provided at the time of the Call.
- Proposal Submission: Potential supervisors must submit a PhD research proposal relevant to one of the specified research areas.
- Review and Shortlisting: An Insight Review Board will assess, and shortlist proposals based on the evaluation criteria outlined below.
- CBI Review and Selection: Shortlisted proposals will be submitted to the CBI for review and funding recommendations.
- Notification and Recruitment: Successful supervisors will be notified and invited to recruit PhD scholars following the host university’s PhD recruitment process, ensuring alignment with the Irish National Framework of Qualifications (Level 10)

### Eligibility Criteria

To be eligible for funding under the CBI PhD Programme, supervisors must meet their institution’s criteria for PhD supervision. The following key principles apply:

- Institutional Affiliation: Applicants must be affiliated with an eligible Irish research institution, such as a university, technological university, or other research-performing organizations recognized by Research Ireland.
- Researcher Status: Applicants must be Principal Investigators, meaning they hold a position that allows them to supervise postgraduate students. This could be a permanent academic position or a contract extending beyond the proposed PhD duration. Early-career Principal Investigators are encouraged to apply as main or co-supervisors.

### Insight Research Ireland Centre for Data Analytics & Central Bank of Ireland

---

- Research Area: Proposed research must align with one of the Call's specified topics.
- Ethical and Legal Compliance: Applicants must ensure their research adheres to ethical and legal requirements, including necessary approvals for studies involving personal data or human participants.
- Residency Requirement: While applicants do not need to be Irish citizens, they must reside in the Republic of Ireland to supervise the student.
- Co-Supervision: Each proposal must include at least one approved Insight Investigator as a supervisor or part of the supervision team. Both interdisciplinary and cross-institutional co-supervision is strongly encouraged to provide broader perspectives on these complex challenges.

### How to apply

Before submitting a proposal, all applicants are encouraged to read this **Call Document** in its entirety.

All proposals must be submitted via the online application form, available [here](https://forms.gle/bMa6pPSdSVCmuDTJ9) [https://forms.gle/bMa6pPSdSVCmuDTJ9]. The link is also accessible via the Insight and Central Bank websites. In addition to completing all mandatory sections within the specified word limits, applicants are required to upload a complete copy of their proposal as a single merged PDF file. This document must include all materials submitted across all sections, and the word count for each response must match exactly what was entered in the corresponding form field.

### Application Form Structure

1. Research Topic:
  - Specify which of the Central Bank's two focus areas this proposal addresses.
2. Lay Abstract\*(Max: 300 words)\*:
  - Provide a summary of the proposal suitable for a non-specialist audience.
3. Research Proposal \*(Max: 500 words)\*
  - Aims, objectives, and central research questions.
  - How existing literature has informed the proposal.
  - Contribution to advancing the state-of-the-art in the field.
4. Research Design & Methodology \*(Max: 500 words)\*
  - Outline the methodologies and approach to be employed.
5. Project Timeline & Risk Management \*(Max: 500 words)\*
  - Key milestones and deliverables.
  - Potential risks and mitigation strategies.
6. Dissemination & Impact Strategy: \*(Max: 500 words)\*
  - Plans for publications, conferences, Education & Public Engagement (EPE), and knowledge exchange.
  - Strategies for measuring the research's impact.

Only fully completed applications<sup>1</sup> received prior to the application deadline will be considered for evaluation. The evaluation consists of a three-stage process:

### Evaluation Process

The evaluation follows a three-stage process:

1. Eligibility Assessment: Administrative compliance check.
2. Technical Evaluation: Assessment based on the evaluation criteria below.
3. Final Decision & Notification: Selection of top-ranked applications for funding.

### Evaluation Rubric

Applications will be assessed based on the following criteria, with the allocated weighting for each section. This list serves as a guideline rather than a strict framework, outlining key aspects assessors will consider.

#### Excellence and Innovation (30%)

- Novelty and ambition in relation to the state-of-the-art in Ireland and beyond.
- Demonstrated understanding of the relevant research landscape.
- Validity and reliability of the proposed concept and approach, including interdisciplinary elements.
- Research track record of the supervisor.

#### Relevance and Impact (30%)

- Alignment with the needs of the CBI and broader research community.
- Relevance to national and EU policies (e.g., National AI Strategy, EU AI Act).
- Contribution to responsible, ethical, and trustworthy AI regulation.
- Potential for influencing policymaking and industry practices.
- Strategies for effective dissemination and stakeholder engagement.
- Ethical data management practices.

#### Quality and Efficiency of Implementation (20%)

- Coherence and effectiveness of the research work plan.
- Quality of the research framework, including clear deliverables, milestones, and a credible breakdown of activities.
- Robust research management and risk mitigation strategies, including scheduling, dependency identification, and monitoring.
- Strength of supervisory management and oversight, including risk management strategies.

---

<sup>1</sup> Fully complete means containing all the relevant information to enable eligibility and technical assessment and as described herein.

### Insight Research Ireland Centre for Data Analytics & Central Bank of Ireland

---

- Eligibility of the supervisor and their institution, ensuring compliance with the programme's requirements.

#### Diversity and Emerging Talent (20%)

- Early-stage supervisors are encouraged to apply as lead or co-supervisors.
- Gender balance among applicant team.
- Evidence of value-added from trans-disciplinary and/or cross-institutional collaboration.

## Conclusions

This call note has outlined the topic areas, eligibility and how to apply, and evaluation process for the 2026 Insight – CBI PhD programme.

All proposals must be submitted via the online application form, available [here](https://forms.gle/bMa6pPSdSVCmuDTJ9) [https://forms.gle/bMa6pPSdSVCmuDTJ9]. The submission link is also accessible via the Insight and Central Bank websites.

Please note the deadline for proposal submission is April 10<sup>th</sup>, 2026.

Any queries should be addressed to the programme email: [cbifellowship@insight-centre.org](mailto:cbifellowship@insight-centre.org)