BANCAS Credit Risk Lab

Student Project Handbook

BANCAS Year 2 · Academic Year 2026–2027

Version 1.1 · 7 October 2026

1. Project challenge and learning objectives

The lab simulates a full credit-risk modelling lifecycle. Your team will receive staged data releases on which you will build an application/behavioral risk model, translate model outputs into a credit decision policy, explain individual decisions, and then test whether the original model and policy remain fit for purpose on a new cohort under changed conditions.

Your objective is to act as a credit-risk team, and build and validate a model, convert calibrated risk estimates into a defensible policy, explain individual decisions, and later determine whether the original model remains reliable on a new cohort.

Prizes

- 1st place: RON 15,000
- 2nd place: RON 10,000
- 3rd place: RON 5,000

Participation in this project will be rewarded with a bonus in the **Banking Operations** course. Winning one of the three prizes will result in an additional bonus, awarded separately within the same course.

2. Data and staged project design

The project uses a synthetic database with customer, account, transaction, product and behavioural information. Data are released in stages so that each milestone reflects a realistic point in the model lifecycle.

2.1 Stage 1 — model development (M1)

- Perform data analysis and quality diagnostics.
- Build a transparent benchmark model, expected to include a logistic scorecard.
- Build non-linear challenger models on the development data.
- Evaluate discrimination and calibration.
- Choose a champion model using statistical, interpretability, stability and operational considerations.
- Translate the champion output into a credit decision policy/risk bands and discuss the expected risk-volume/economic trade-off.
- Make the complete workflow reproducible from the repository.

2.2 Stage 2 — Client decision challenge

After M1 is frozen, you will analyse a specified case using your frozen M1 model/policy. You must reconstruct the decision, identify the main local drivers, provide understandable reason codes and distinguish predictive explanations from causal statements.

2.3 Stage 3 — Out-of-time (OOT) monitoring and redevelopment

A later release introduces a new/out-of-time cohort and permitted new information

- Compare OOT discrimination and calibration with M1.
- Test whether the original threshold/policy remains suitable under the new population and economics.
- Evaluate permitted new variables and challenger/redevelopment options.
- Make a final recommendation: keep, recalibrate, change threshold/policy, segment, or redevelop.

3. Milestones and deliverables

Milestone	Weight	Required tasks and outcomes
M1 – Model Development (Stage 1)	25%	Data analysis and quality assessment; development of a benchmark logistic scorecard and challenger models; evaluation of classification and calibration performance; champion model selection (**Stage 1 Model, M1**); definition of the credit decision policy and its economic implications; full reproducibility of the analytical flow.
Client Decision Challenge (Stage 2)	10%	Business interpretation of a decision to approve or decline a credit application based on Stage 1 model (**M1**) and policy: reconstruction of the decision for the specific case, including fairness and governance considerations.
M2 – OOT Monitoring and Redevelopment (Stage 3)	25%	Assessment of the Stage 1 model (**M1**) on the OOT sample (OOT monitoring); evaluation of classification and calibration performance; model stability (Population Stability Index); factors of performance deterioration; recommendation on whether to keep or redevelop the model; policy impact and recommendation.
Final Committee	40%	Presentation of the financial recommendation of the end-to-end model lifecycle: covering model development, credit decision-making, monitoring and redevelopment (in English).

4. Calendar

Date / period	Stage	Expected activity
07–13 Oct 2026	Team registration	Each student submits the linked form and chooses one of eight teams: Tlaloc, Agamemnon, Ajax, Juno, Barbarossa, Xerxes, Dante or Hecate. (see section 6.1) Organizers confirm placement; maximum 3 students per team.
13–18 Oct	Week 1	Onboarding, project brief, data architecture, reproducibility setup, initial data quality
19–22 Oct	Week 2	Data analysis, application-time feature inventory, target/outcome logic, baseline design. **Mentorships**: exact session details will be shared in advance.
2–8 Nov	Week 4	Benchmark scorecard and modelling pipeline; cross-validation on development data.
9–15 Nov	Week 5 / M1	Model development, calibration, policy/economics and M1 freeze. M1 due by the end of 15 Nov. **Mentorships**: exact session details will be shared in advance.
16–22 Nov	Client decision challenge week	Local decision reconstruction, explainability, fairness/governance review.

Date / period	Stage	Expected activity
23–29 Nov	Week 7	M2 release: incumbent monitoring on new/OOT cohort; drift and stability analysis; permitted enrichment variables.
2–4 Dec	Week 8 / M2	Redevelopment/challenger comparison; recalibration or policy changes where justified. M2 freeze by end of 6 Dec. **Mentorships:** exact session details will be shared in advance.
7–11 Dec	Final week	Final committee presentation and defence; FINAL freeze before the assigned committee slot.

5. Mentoring workflow

Mentoring is organised inside GitHub so that questions and decisions remain traceable and no email submission process is required.

Mentorships are planned during 19–23 October, 9–13 November and 1–4 December. The organisers will share the exact session date and time in advance.

1. **At least 24 hours before** a scheduled mentoring session, open a Mentoring Pre-read issue in your private team repository.
2. Include a brief status update, decisions since the previous session, links to the relevant committed files and no more than three priority questions.
3. Use the session to test assumptions and decisions; the mentor will not edit your code or choose the model for you.
4. After the session, record conclusions, owners and next actions in the same issue.
5. For a clarification that applies to all teams, use the General Project Question issue in project-hub.

6. GitHub workflow and repository rules

Component	Configuration	What it means for you
Organisation	bancas-credit-risk-lab	Private academic workspace.
project-hub	Private shared repo	Student instructions, project brief/calendar, FAQ, shared questions and registration-form link.
team-template	Private template	Source for each team repo; notebooks/, src/, reports/, predictions/, submissions/, data/README.md and validation tooling.
instructors	Private teacher-only repo	Administration, grading, private evaluation materials and mentor allocations.
Organisation base access	No permission	Access is granted explicitly; students must not be able to browse other team repositories.

Component	Configuration	What it means for you
Team repositories	Eight private repositories provisioned from team-template.	Students: Write access only to their confirmed team repository. Mentors&Instructors: Read access across all eight private team repositories for project review.
• Tlaloc	Open private repository	Private workspace; access is granted after team confirmation.
• Agamemnon	Open private repository	Private workspace; access is granted after team confirmation.
• Ajax	Open private repository	Private workspace; access is granted after team confirmation.
• Juno	Open private repository	Private workspace; access is granted after team confirmation.
• Barbarossa	Open private repository	Private workspace; access is granted after team confirmation.
• Xerxes	Open private repository	Private workspace; access is granted after team confirmation.
• Dante	Open private repository	Private workspace; access is granted after team confirmation.
• Hecate	Open private repository	Private workspace; access is granted after team confirmation.
Milestone tags	M1-FROZEN; XAI-FROZEN; M2-FROZEN; FINAL-FROZEN	Tags point to immutable submitted states; the recorded full commit SHA is authoritative.

6.1 Getting access

1. Create a personal GitHub account.
2. Submit the team registration form with your full name, GitHub username and one chosen team. Each student submits individually. Do not enter an email address, phone number or password. The form is accessible to anyone with its link; responses are visible only to the organisers. A team choice is a request, not a reserved place.
3. Each team can have up to 3 students. Organisers review team choices and available places, confirm your team, and send you an invitation to the GitHub organisation. Accept the invitation.
4. After your team is confirmed and you have accepted the organisation invitation, open only your confirmed team repository using the links below. The repository is private and the organisers grant access only to your confirmed team repository. Do not create repositories yourself or share another team's private work.

- Tlaloc
- Agamemnon
- Ajax
- Juno
- Barbarossa
- Xerxes
- Dante
- Hecate

5. Do not create your own project repository. Report immediately if you can see another team's private repository.

6.2 Where work belongs

Path	Use
notebooks/	Exploration and modelling notebooks.
src/	Reusable code.
reports/	Milestone reports and one-page decision summaries.
predictions/	Prediction CSVs when required.
submissions/	Milestone manifest JSON files.
data/	Follow data/README.md. Do not commit data unless explicitly authorised.
README.md	Reproduction entry point, dependencies, execution order, seeds/runtime assumptions.

Commit regularly with short, meaningful messages. Pull the latest changes before editing shared notebooks/files. Never commit credentials, real customer information or any material that was not explicitly released for the project.

7. Submission and freeze procedure

A milestone is not submitted by uploading a report alone. Your frozen repository state is the academic evidence.

1. Commit and push all intended deliverables.
2. Confirm that git status shows no intended uncommitted files.
3. Create and push the milestone tag: M1-FROZEN, CDC-FROZEN, M2-FROZEN or FINAL-FROZEN.
4. Copy the full 40-character commit SHA.
5. Check the GitHub Actions validation run.
6. Open a Milestone Submission issue. Include the tag, full SHA, commit-based links to deliverables, validation run and one-page decision summary.
7. Both the submitted commit and the submission issue must meet the published deadline.
8. Do not delete, move or silently replace a submission tag. If a correction is needed, request approval in the existing submission issue.

8. Reproducibility, validation and academic standards

- Use only information that would have been available at the stated decision/observation time.
- Do not use test or OOT information to select features, tune hyperparameters or choose a model.
- Keep missing outcomes missing unless the assignment explicitly provides a valid label; never create labels because they are convenient for modelling.
- Document exclusions, imputations, transformations, calibration choices and policy thresholds.
- Your model must be reproducible from the frozen commit with the stated dependencies and execution order.
- You are responsible for understanding and defending all code and analytical choices, including work produced with AI assistance.
- Use discrimination, calibration and stability evidence together. A model can rank well and still be unsuitable as a probability or decision model.