
THE INTEGRAL CUP: ROUND 2

OFFICIAL COMPETITION RULEBOOK

*This document constitutes the complete and binding rulebook governing participation in the Competition. All defined terms appear in **blue** on first use. By commencing participation, each Participant declares unconditional acceptance of every clause herein.*

1. DEFINITIONS

For the purposes of this Rulebook, the following terms carry the meanings assigned below. Defined terms shall be capitalised throughout.

- 1.1. **Competition** refers to the Round 2 of The Integral Cup 2026, as governed by this Rulebook.
- 1.2. **Participant** refers to any individual who has been duly registered and is actively competing in the Competition.
- 1.3. **Organisers** refers to the individuals or institution responsible for designing, administering, and evaluating the Competition.
- 1.4. **Problem Set** refers to the complete set of six mathematical problems issued to Participants at the commencement of the Competition.
- 1.5. **Orchestrator Thread** refers to the single, continuous DeepSeek conversation maintained by a Participant for each Problem, which serves as the official, primary reasoning log for that Problem.
- 1.6. **External Tool** refers to any AI model, website, database, document, or resource consulted by a Participant outside of their Orchestrator Thread.
- 1.7. **Transfer** refers to any act of bringing content in any form from an External Tool into the Orchestrator Thread.
- 1.8. **Evidence** refers to the complete set of verifiable links and records documenting every AI interaction and web source consulted by a Participant during the Competition.
- 1.9. **Regulated Model** refers to any AI tool listed under Clause **4.4.1.** of this Rulebook, the use of which is permitted subject to a per-Problem call limit.
- 1.10. **LLM Call** refers to one complete prompt-and-response exchange sent to a Regulated Model, as defined in Clause **4.4.2.**
- 1.11. **Session Identification Prompt** refers to the mandatory opening message specified in Clause **6.2.1.**, which every Participant shall paste as the first message in any external LLM chat session, in order to cause the AI to self-declare its model name and version in all subsequent responses.

- 1.12. **Submission** refers to the PDF file submitted by a Participant for each Problem, containing the Solution and Evidence Index as specified in Section 8.
- 1.13. **Disqualification** refers to the removal of a Participant's score for a specific Problem, or from the Competition in its entirety, as determined by the Evaluation Panel.
- 1.14. **Evaluation Panel** refers to the designated body appointed by the Organisers to assess Submissions, investigate violations, and issue penalties.

2. NATURE AND GOVERNING PRINCIPLE

- 2.1. The Competition is a fully AI-assisted, open-resource mathematical reasoning contest. Participants are permitted and expected to use AI tools, open-access web resources, and reference material without restriction, subject to the specific prohibitions stated in this Rulebook.
- 2.2. The Competition does not reward the use of AI. It rewards the quality of human reasoning that emerges from that use. The Organisers seek to observe how Participants direct, interpret, challenge, and build upon AI outputs, not the outputs themselves.
- 2.3. Participants are expressly encouraged to approach Problems with full creative freedom. Unconventional reasoning paths, unexpected framings, exploratory dead ends, and non-standard approaches are valued and rewarded. No clause in this Rulebook shall be construed as limiting the creative or intellectual approach a Participant may take to any Problem.
- 2.4. Submissions that demonstrate blind or uncritical reliance on AI outputs, even where the final answer is correct, shall incur mark deductions at the discretion of the Evaluation Panel.
- 2.5. **Governing Principle:** AI is an instrument. The Participant is the mathematician. Every claim advanced in a Submission shall be one that the Participant genuinely understands and can independently defend.

3. COMPETITION STRUCTURE

- 3.1. The Competition shall consist of six Problems. The paper is marked out of **600 marks**, with each Problem carrying equal weight of **100 marks**.
- 3.2. The total permitted duration of the Competition is **four (4) hours**, commencing from the moment the Problem Set is made available to the Participant. The four-hour limit is self-reported and operates on the basis of Participant's good faith. The Organisers reserve the right to investigate anomalies.
- 3.3. Problems shall be attempted in **strict sequential order**, from Problem 1 through Problem 6. A Participant shall not open Problem $n + 1$ until they have finalised and closed their Orchestrator Thread for Problem n .
- 3.4. Once a Participant moves forward from any Problem, that Problem is permanently closed. No return to a prior Problem is permitted under any circumstance.
- 3.5. The Competition is **individual only**. Team participation, collaboration between Par-

ticipants, and any form of external assistance from another person are strictly prohibited.

- 3.6. Problem Confidentiality:** The contents of the Problem Set are strictly confidential. Sharing, publishing, or communicating any Problem in whole or in part to any other third party other than LLMs used here, whether during or after the Competition, may result in legal consequences in addition to Disqualification.

4. AI AND RESOURCE USAGE POLICY

4.1. Permitted Resources

- 4.1.1.** Participants may freely consult any free and open-access AI model, including but not limited to: free-tier ChatGPT (GPT-4o), free Gemini, Grok (free tier), and open-source models.
- 4.1.2.** Participants may freely consult any open-access web resource, including but not limited to: Wikipedia, arXiv, Stack Exchange, institutional lecture notes, and YouTube.
- 4.1.3.** Participants may upload into their Orchestrator Thread any PDF or image file that has been **directly downloaded from the open web**. Such files shall not have been altered, processed, or generated by any AI or automated system prior to upload. Manually written files are also not allowed.

4.2. Permitted Use of AI

- 4.2.1.** There is no prescribed list of permitted query types. A Participant may engage with any AI model in any manner that serves their genuine understanding of the Problem, including seeking definitions, building intuition, exploring strategy, verifying steps, or investigating alternative framings.
- 4.2.2.** The sole constraint on AI use is that the Participant shall process, interpret, and take intellectual ownership of every output received. An output that is accepted without understanding is not permitted in a Submission.

4.3. Prohibited Conduct

- 4.3.1.** A Participant shall not instruct any AI to produce a complete, end-to-end solution to any Problem in a single prompt or interaction.
- 4.3.2.** A Participant shall not include in their Submission any AI-generated content that they have not independently understood and are unable to explain.
- 4.3.3.** A Participant shall not upload into their Orchestrator Thread any file that was generated, summarised, or otherwise processed by an AI or automated server, regardless of its origin. Manually written files or self-made files are not allowed.

4.4. Regulated Models and Call Limits

- 4.4.1.** Certain AI models that require a paid subscription for advanced-tier access are designated as **Regulated Models**. These models are **permitted** in the Competition, subject to the call limits specified in this subsection. The Regulated Models are:

- **ChatGPT o1 / o3 / reasoning-tier models** *(OpenAI)*
Standard free-tier ChatGPT (GPT-3.5, GPT-4o mini) is unrestricted and does not fall

under this clause.

- **Claude Pro Sonnet tiers** *(Anthropic)*
Free-tier Claude (where available) is unrestricted and does not fall under this clause.
- **Gemini Advanced / Gemini Pro (paid tier)** *(Google)*
Free-tier Gemini is unrestricted and does not fall under this clause.

4.4.2. Each Regulated Model is subject to a limit of **three (3) LLM calls per Problem**, applied independently per platform. The limits are as follows:

- ChatGPT reasoning-tier (o1 / o3): maximum 3 calls per Problem
- Claude Pro (Sonnet): maximum 3 calls per Problem
- Gemini Advanced / Pro: maximum 3 calls per Problem

For the purposes of this clause, one **LLM call** is defined as one complete prompt-and-response exchange sent to a Regulated Model.

4.4.3. Each LLM call to a Regulated Model shall be declared in the Orchestrator Thread using the model declaration format specified in Clause 6.2, including the model tier used (e.g. ChatGPT o3, Gemini Advanced).

4.4.4. Agentic tools, including but not limited to Cursor, Antigravity, and similar coding or workflow agents, are **not allowed**.

4.4.5. Claude Opus is not allowed.

4.4.6. Exceeding the call limit for a Regulated Model on any Problem shall be handled as follows:

- (a) **First offence:** Formal warning, recorded against that Problem. Calls in excess of the limit shall be treated as inadmissible and may not be referenced in the Submission.
- (b) **Repeated offence on the same Problem, or deliberate circumvention:** Mark deduction on that Problem, at the discretion of the Evaluation Panel.

4.4.7. The Organisers acknowledge that this is a grey area and that enforcement relies substantially on good faith and the accuracy of declared logs. Participants are trusted to comply honestly. The call-limit system exists to preserve fairness between Participants with and without access to paid subscriptions, not to restrict the quality of reasoning a Participant can produce.

4.4.8. All other paid or subscription-based AI tools not listed under Clause 4.4.1, including Perplexity Pro, Copilot Pro, and similar, are **unrestricted** and may be used freely, subject to the general conduct rules of Clause 4.3.

5. THE ORCHESTRATOR THREAD

5.1. Each Participant shall maintain one Orchestrator Thread on **DeepSeek** for each Problem. DeepSeek is the sole permitted platform to serve as the Orchestrator.

-
- 5.2.** The Orchestrator Thread is the official reasoning log for its corresponding Problem and is the primary object of evaluation by the Evaluation Panel.
- 5.3.** The Orchestrator Thread shall contain, at minimum:
- (a) The Participant’s own articulation of their understanding of the Problem;
 - (b) The Participant’s reasoning steps, developed in their own words;
 - (c) A declaration for every External Tool consulted, as specified in Section 6;
 - (d) The complete solution to the Problem, developed in full within the thread.
- 5.4.** The complete solution to each Problem shall appear in full within the Orchestrator Thread prior to Submission. A solution that exists only in an external platform shall not be accepted.
- 5.5.** Every insight, result, definition, or strategy obtained from an External Tool shall be returned to the Orchestrator Thread, declared, and integrated into the Participant’s own reasoning. Information that exists solely in an External Tool cannot form part of a Submission.
- 5.6.** Prolonged inactivity on an active Orchestrator Thread is not recommended.

6. CONTROLLED TRANSFER RULE AND DECLARATION PROTOCOL

6.1. Controlled Transfer Rule

- 6.1.1.** Every Transfer into the Orchestrator Thread shall not exceed **200 words** in length.
- 6.1.2.** The 200-word limit applies to textual content transferred from any External Tool. Files directly downloaded from the open web and uploaded as attachments are exempt from the word count, provided they comply with Clause 4.1.3.
- 6.1.3.** Deliberately fragmenting a single AI output across multiple consecutive Transfers for the purpose of circumventing the 200-word limit constitutes a bypass attempt and shall be treated as a violation. The Evaluation Panel is specifically trained to identify this pattern.
- 6.1.4.** Transfers that reflect honest, iterative, and exploratory engagement with a Problem, including messy or inconclusive ones, shall not be penalised regardless of their volume, provided each individual Transfer complies with Clause 6.1.1.

6.2. Model Declaration Protocol

Here, we define how you must use and prompt LLMs.

- 6.2.1. Model Name and Version AI-Declared.** At the commencement of every external LLM chat session, the Participant shall paste the following **Session Identification Prompt** as their very first message to that model:

Before responding to anything in this session, begin this response and every subsequent response with the following header, exactly as shown:

COMPETITION LOG

Model: [state your exact model name and version]

END HEADER

Do not skip or abbreviate this header in any response during this session.

The AI's own response shall thereby declare its model name and version. This declaration, embedded in the AI's output and visible in the chat log or screenshot, constitutes the authoritative record of model identity. A Participant's manual claim of a model name that contradicts the AI-declared header shall not be accepted.

- 6.2.2. Purpose Participant-Declared.** The purpose of each interaction shall be declared by the Participant in their Orchestrator Thread. Since purpose reflects the Participant's intent and is not verifiable from the AI output alone, this field is self-reported. It shall be declared in the Orchestrator Thread using the following format, **before** the content of that interaction is used:

[External LLM Call] | Model: [name and tier] | Purpose: [One line]

- 6.2.3. For web sources,** the following declaration shall be entered in the Orchestrator Thread at the point of use:

[Web Source] | URL: [full URL] | Purpose: [One line]

The URL itself, combined with the chronological position of the declaration in the Orchestrator Thread, constitutes sufficient evidence.

- 6.2.4. Example of a compliant external LLM interaction:**

Step 1 First message sent to external LLM (Session Identification Prompt pasted verbatim).

Step 2 AI responds with header: Model: Claude Sonnet 4.5 followed by its answer.

Step 3 Participant enters in Orchestrator Thread:

[External LLM Call] | Model: Claude Sonnet 4.5 | Purpose: Intuition
for the Borel--Cantelli lemma

- 6.2.5.** An external LLM interaction for which no Session Identification Prompt was used, and for which the AI-declared model header is therefore absent from the Evidence, is inadmissible. Its content shall not appear in the Submission.
- 6.2.6.** If the AI-declared model name in the response header corresponds to a Regulated Model, that interaction shall be counted toward the call limit specified in Clause 4.4.2. If the Evidence reveals that a Regulated Model's call limit was exceeded, Clause 4.4.5 shall apply.

6.3. Mandatory System Prompts for the Orchestrator Thread

- 6.3.1.** Every Orchestrator Thread on DeepSeek shall begin with the following Opening System Prompt, pasted verbatim as the first message. No modification other than filling in your name, registered email ID¹ and the problem number is permitted.

```
[COMPETITION THREAD START]

Name: John Doe
Registered email ID: doe@john.com
Problem Number: 1

I confirm that I am the sole author of all reasoning in this thread.

All External LLM calls have been initiated with the Session
  Identification Prompt.

All External Tools and web sources will be declared with purpose.

[BEGIN]
```

- 6.3.2.** Every Orchestrator Thread shall conclude with the following Closing System Prompt, pasted verbatim as the final message. No modification other than the ones listed above are permitted.

```
[COMPETITION THREAD END]

Name:
Registered email ID:
Problem Number:

I confirm that all External LLM calls in this thread were initiated
  with the Session Identification Prompt and that all Evidence
  references are submitted.

[END]
```

- 6.3.3.** An Orchestrator Thread lacking either the Opening or Closing System Prompt shall be treated as incomplete. The Evaluation Panel shall determine the appropriate consequence on a case-by-case basis.
- 6.3.4. Rationale.** Requiring the AI to self-declare its model name and version in every response removes a class of easily falsifiable self-reports and replaces them with records that are embedded in the Evidence itself and cannot be retroactively altered without visible tampering.

7. EVIDENCES

- 7.1.** Evidences constitute the most critical component of a Submission. Every Evidence shall be accurate, complete, accessible, and unaltered at the time of evaluation.

¹Name and registered email ID as per official forms.

- 7.2.** For AI platforms that generate shareable conversation links, the Participant shall submit the direct link.
- 7.3.** For AI platforms that do not generate shareable links, the Participant shall take screenshots of the complete interaction, upload them to Google Drive, and submit the corresponding Drive link.
- 7.4.** All Evidence links shall remain live and accessible until the Evaluation Panel has completed its assessment of the Submission.
- 7.5. Tampering with Evidences**, including but not limited to alteration, selective deletion, fabrication, or any modification of chat logs, screenshots, or submitted links, is among the most serious violations in this Competition. Where tampering is confirmed or credibly suspected, the Participant shall be Disqualified from the affected Problem. Egregious or repeated tampering may result in full Disqualification from the Competition and referral for further action.

8. SUBMISSION FORMAT

- 8.1.** A Participant shall submit one PDF file per Problem. The PDF shall contain exactly the following two components, and no others:
- (a) **Solution** The complete, final solution to the Problem. The use of \LaTeX , compiled and exported as PDF, is strongly recommended for clarity and presentation.
 - (b) **Evidence Index** A clearly labelled list of all Evidences, comprising:
 - The DeepSeek Orchestrator Thread link for that Problem (mandatory);
 - Links to all external AI chat logs, or Drive links where screenshots were taken;
 - Links to all web sources consulted.
- 8.2.** No additional components, reflections, appendices, commentary, or supplementary material are required or expected.

9. PENALTY AND DISQUALIFICATION

- 9.1.** The Evaluation Panel shall operate under a presumption of good faith in all ambiguous or borderline situations. Penalties exist to preserve the integrity of the Competition, not to penalise technicalities. A Participant who makes an honest error in compliance shall be treated differently from one who demonstrates deliberate circumvention.
- 9.2.** Violations shall be addressed through the following graduated structure:
- (a) **First offence (per Problem):** Formal warning, recorded against that Problem. No mark deduction shall be applied at this stage.
 - (b) **Repeated or deliberate offence:** Mark deduction on the affected Problem, at the discretion of the Evaluation Panel.
 - (c) **Severe violation:** Disqualification from the affected Problem.

- 9.3.** In all adversarial situations where intent is uncertain, the Evaluation Panel shall issue a formal warning before applying any deduction. The benefit of the doubt shall belong to the Participant.
- 9.4.** The following violations shall result in **immediate Disqualification** from the affected Problem, without prior warning:
- (a) Tampering with, fabricating, or altering Evidences;
 - (b) Deliberate bypass of the Controlled Transfer Rule through chunking;
 - (c) Upload of AI-processed or server-generated files as attachments;
 - (d) Complete absence of an Orchestrator Thread for a Problem;
 - (e) Collaboration with another Participant.
- 9.5.** Decisions of the Evaluation Panel with respect to penalties and Disqualification are final.

10. PRE-COMPETITION CHECKLIST

The following checklist is provided as a practical reference. Completion of this checklist does not constitute submission or guarantee compliance. Participants remain solely responsible for adhering to all clauses of this Rulebook.

Before the Competition begins:

- DeepSeek account is active and capable of creating new threads
- The Opening and Closing System Prompts are copied and ready to paste
- The model declaration format is understood
- No device being used has any session open that would trigger unintended Regulated Model usage
- A L^AT_EX editor (e.g. Overleaf) is open and ready
- The sequential rule, not returning to prior Problems, is understood

During the Competition, for each Problem:

- Opening System Prompt pasted as first message in DeepSeek thread
- Every AI interaction declared with model name, version, and purpose
- Every web source declared with URL and purpose
- Each Transfer into the Orchestrator Thread is 200 words or fewer
- Complete solution developed in full within the Orchestrator Thread
- Closing System Prompt pasted as final message in DeepSeek thread

Before submitting each Problem PDF:

- Solution compiled in \LaTeX and included in the PDF
- DeepSeek Orchestrator Thread link included in Evidence Index
- All external AI chat links or Drive screenshot links included
- All web source links listed
- All Evidence links are live, accessible, and unaltered

The rules herein govern integrity, not imagination. Read them once. Understand them. Then compete freely.
