


When several AI models work on the same task, inconsistent inputs create inconsistent answers. One model may focus on the financials, another may assume a different audience, and a third may fill gaps with confident-sounding guesses. Comparing those outputs is difficult because you are not comparing like with like.

The fix is a shared brief: a clear, reusable record of the task, its audience constraints, the source material, and the defined output. Give every model the same brief and the same evidence, and differences in their responses become more useful. You can see where they agree, where they interpret the evidence differently, and what needs human judgment.

This guide explains what to put in the brief, how to use it in a multi-model AI chat in one shared thread, and when to run models in parallel or in sequence. It also uses a hypothetical \$42M acquisition example involving Skybridge to make the approach concrete.

Why a shared brief matters

A prompt is often treated as a one-off instruction. A brief is more like a small operating system for the work: it tells each model what problem to solve, what evidence it may use, what assumptions to avoid, and how to present its answer.



Without that structure, you can accidentally introduce variation before the models even begin. Perhaps one gets a long background explanation while another sees a short prompt. Perhaps one receives a spreadsheet and another receives only a summary. If their conclusions differ, you will not know whether the difference came from their reasoning or their inputs.

A shared brief creates a fair comparison. It does not make models identical, nor should it. Instead, it gives them the same starting conditions so that their different interpretations can help you make a better decision.

The essential parts of a model-ready brief

1. The decision or job to be done

Start with the actual decision the work is meant to support. Avoid broad requests such as “analyze this deal” or “review this plan.” Name the question, the decision-maker, and what a useful answer will help them do.

For example: “Help the investment committee identify the most important diligence questions before deciding whether to proceed to the next stage.” That is more actionable than “summarize the acquisition.”

2. Audience constraints

Specify who will read or use the answer and what that audience needs. Audience constraints can include seniority, subject-matter expertise, time available, risk tolerance, and whether the document is for internal or external use.

- **Reader:** investment committee, operating team, customer, or general audience
- **Expertise:** specialist, informed generalist, or new to the topic
- **Use:** decision, discussion, briefing, or publication
- **Tone:** neutral, persuasive, cautious, or exploratory
- **Constraints:** confidentiality, regulatory sensitivity, or limited reading time


Be specific. “Write for executives” leaves room for interpretation. “Write for a five-person investment committee that knows M&A but has not reviewed the underlying materials; keep the summary to two minutes of reading” gives each model a clearer target.

3. Source material and evidence rules

List the source material available to the models and explain how to treat it. Name files or sections where possible, identify the date or version, and distinguish primary evidence from summaries or commentary.

- What documents are in scope?
- Which source is authoritative if two sources conflict?
- May the model use general background knowledge, or must it stick to the supplied material?
- How should it handle missing, ambiguous, or conflicting information?
- Should every material claim be tied to a page, section, or source?

A useful instruction is: “Use only the supplied materials for deal-specific facts. Do not infer that an unreported metric is zero. Mark unavailable information as ‘not provided’ and identify what source would resolve the gap.” That makes uncertainty visible instead of inviting the model to fill it in.



4. Scope, assumptions, and exclusions

Define the edges of the task. State what to analyze, what not to analyze, and which assumptions are permitted. If a model must make an assumption to proceed, tell it to label the assumption and explain how it affects the conclusion.

For instance, you might ask for a commercial diligence review but exclude legal and tax advice. You might request a preliminary risk assessment while making clear that the available materials are incomplete. These boundaries help prevent a polished answer from being mistaken for a complete investigation.

5. A defined output

Tell each model exactly what to return. A defined output makes responses easier to compare and reduces the chance that one model produces a narrative while another produces a list of unrelated observations.

Specify the structure, length, level of detail, and any required labels. For example:

- Executive summary: no more than 150 words
- Top five decision-relevant findings, ranked by importance
- For each finding: evidence, implication, confidence, and open question
- Separate facts from assumptions and recommendations
- End with three diligence requests that would most change the decision

Also say what a good answer should not do. "Do not give a generic M&A checklist" may be as important as "rank the findings."

A practical brief template

Use this template as a starting point. Keep the core brief stable when sending the task to different models; add model-specific instructions only when they are essential.

Brief section What to include Task The decision or question the work must support Audience constraints Reader, expertise, use case, tone, and length expectations Source material Documents, versions, evidence rules, and

citation expectations Scope Included topics, exclusions, and boundaries Assumptions Allowed assumptions and how to label uncertainty Defined output Required structure, ranking, word limits, and ending Review criteria What makes the answer useful, accurate, and decision-relevant

Example: a \$42M acquisition involving Skybridge

Imagine a team evaluating a hypothetical \$42M acquisition involving a company called Skybridge. The team wants an early view of the deal's key commercial risks. The brief should not imply facts that are not in the materials; it should tell every model how to analyze the same evidence.

Task: Identify the five commercial diligence questions most likely to affect whether the committee proceeds to the next stage of evaluating the \$42M Skybridge acquisition.

Audience constraints: Write for an investment committee familiar with acquisitions but new to the supplied Skybridge materials. Use concise, neutral language. Prioritize decision relevance over background explanation.

Source material: Use only the provided management presentation, customer summary, and financial workbook. Refer to the source and page, tab, or section for each material deal-specific claim. If evidence is absent or sources conflict, flag the issue rather than resolving it by guesswork.

Scope: Focus on revenue durability, customer concentration, growth assumptions, and integration dependencies. Do not provide legal, tax, or valuation advice. Do not treat the \$42M price as evidence that the deal is attractive or unattractive.

Defined output: Return a 100-word executive summary, followed by five ranked diligence questions. For each question, include the evidence, why it could matter to the decision, confidence level, and the specific information needed to resolve it. End with the three documents or analyses that should be requested next.

This example gives models a common frame without telling them what conclusion to reach. One might identify a concern in customer concentration; another might notice that the growth case depends on an unsupported assumption. Both are useful if each ties its point to the same source material and makes uncertainty explicit.

Use one shared thread to preserve the inputs

In a multi-model AI chat in one shared thread, keep the brief and source material together wherever the product allows it. That creates a visible record of the task, the evidence provided, and the instructions used. It also makes it easier for teammates to review the process and spot when a model received an updated file or an extra instruction.

Before comparing responses, check that the models actually received the same materials. "Same prompt" is not enough if one model cannot access an attachment, sees a different document version, or has additional context from earlier in the conversation. A shared thread can reduce this risk, but only if the workflow preserves the same inputs in practice.

For repeatable work, label the brief and source set with a version or date. If you revise the brief, note what changed and rerun the comparison when the change could affect the result. This prevents an output from looking comparable when its instructions were not.

Choose parallel or sequential workflows deliberately

Parallel workflows: independent first opinions

In a parallel workflow, several models receive the same brief and source material at roughly the same time. This is useful when you want independent perspectives, a broad set of risks, or a quick comparison of interpretations.

Ask each model to work independently before it sees the others' answers. If models read one another's conclusions first, they may converge on an early framing instead of surfacing distinct evidence. For the Skybridge example, parallel review can show whether multiple models independently flag the same diligence gap, while also revealing issues noticed by only one.

Sequential workflows: structured handoffs

In a sequential workflow, one step builds on another. A model might extract evidence from source documents, a second might analyze the extracted evidence, and a third might draft a committee memo from the analysis.

Sequence is useful when the task has clear stages or when a later step needs a standardized intermediate output. It also creates a risk: errors or assumptions in an early step can travel downstream. Preserve the source references, label uncertain extractions, and review intermediate work before asking the next model to rely on it.

A simple rule: use parallel work to generate independent analysis; use sequential work to move through a controlled process. For high-stakes decisions, you can combine them: gather independent first-pass analyses, reconcile disagreements, then use a final drafting step to produce the defined output.

Use orchestration modes for structured work

Orchestration modes describe how multiple models or steps are coordinated. The mode should match the job, not simply the number of models available.

- **Independent review:** Models answer the same brief separately. Best for surfacing distinct readings of the evidence.
- **Specialist review:** Models receive the same core brief but focus on assigned lenses, such as customers, financial assumptions, or operations. Best when the task has clearly separable workstreams.
- **Review and challenge:** One model drafts an analysis and another checks its evidence, logic, and missing questions. Best for finding gaps, but not a substitute for an independent first pass.
- **Draft and refine:** A model produces a structured draft, then another improves clarity against the same criteria. Best for writing tasks after the underlying analysis has been reviewed.

When you assign different roles, preserve a shared core brief and make each role's additional instruction explicit. Otherwise, role-specific prompts can become hidden differences that make results hard to compare.

Treat disagreement as a signal

Disagreement between models is not automatically a sign that one is wrong. It can be a signal to inspect the evidence, the wording of the task, or the assumptions behind a conclusion.

Sort disagreements into useful categories:

- **Evidence disagreement:** Models cite or interpret different parts of the source material.
- **Assumption disagreement:** They make different assumptions where the materials are silent.
- **Priority disagreement:** They notice similar issues but rank their importance differently.
- **Fact error:** A model misreads, invents, or misattributes information.

For each substantive disagreement, ask: What exact evidence supports each view? Is the difference caused by an unclear definition or missing data? What additional information would settle the issue? In the Skybridge example, if one model ranks customer concentration as the top concern and another ranks growth assumptions first, the team can test both against the workbook and customer materials instead of choosing whichever answer sounds more confident.

Disagreement becomes especially useful when the brief requires citations, confidence levels, and open questions. Those fields make it easier to distinguish a genuine interpretive difference from a factual mistake—and turn model output into a focused next step for human review.

Make the brief repeatable

Before launching a multi-model review, run a short checklist:

1. Does every model have the same task, audience constraints, and source material?
2. Are the scope, exclusions, and assumptions explicit?
[multi AI for grant writing](#)
3. Is the defined output specific enough to compare side by side?
4. Can reviewers trace important claims back to the evidence?
5. Have you chosen parallel, sequential, or specialist orchestration intentionally?
6. Is there a human reviewer responsible for resolving uncertainty and making the decision?

A strong brief will not guarantee a correct answer. It will make the work more consistent, the differences more interpretable, and the gaps easier to investigate. When every model starts with the same inputs, the goal is not to make them agree. It is to create better conditions for finding what matters, testing what is uncertain, and making a more informed decision.