

PRACTICE B

Paper to Talk, *Then to Referee*

20 minutes, hands on — build a presentation and a referee report.

Jiaqi Shao · NUS Business School



Two artifacts, one paper — present it fairly, then referee it fairly.

■ The Next 20 Minutes

You verified the paper — now scale up to two finished deliverables.

- Goal: a faithful *talk* (`present.tex`) and a fair *referee report* (`critic.tex`).
- Levers from the lecture: plan mode, a *skill*, and a *subagent*.
- Your instructor's demo shows the target — yours need not match slide-for-slide, only fact-for-fact.



■ 1 Plan Before You Draft (4 min)

Plan Mode is a read-only gate — see the outline before any file is written. Press Shift+Tab to enter Plan Mode, then:

Plan a faithful conference talk from @RCT_paper.md using the paper-to-slides skill. Show me the slide outline before drafting.

- Check the arc: motivation → question → design → main result → mechanism → limits.
- Adjust in words (“add a slide on heterogeneity”), then approve to let it write.

■ 2 Build the Talk (8 min)

A skill is packaged expertise — it already knows the house style and the faithfulness rule.

Use the paper-to-slides skill to build `my_slides/present.tex`, then compile it with `xelatex`.

- It first writes a `paper_brief.md` — every number with its table.
- Then drafts, compiles, and re-checks each number.

```
> open present.pdf  
# skim it against  
# the paper
```

If a slide shows a number with no table, that is the rule failing — ask it to source or cut it.

■ 3 Referee It (8 min)


A subagent has its own brief — here, the ethos of a fair referee.

Use the referee subagent to write `my_slides/critic.tex`: recommendation up front, credit first, then majors --- each citing a table and a concrete ask.

- A *fair* referee criticises only what the paper does, and never re-demands a check it already ran.
- Each major carries a “*what would change my mind*” — a runnable request, not a complaint.

■ Do Both at Once (optional)

The talk and the referee report are independent — so run them in parallel.



Use two parallel subagents --- one uses the paper-to-slides skill to build the talk, one acts as referee. Report back when both compile.

■ Checkpoint

Two decks that a co-author and an editor would each recognise.

■ What good looks like

Talk: ~20 slides, motivation to takeaways, every statistic naming its table.

Referee: recommendation up front, credit before criticism, ~5 majors — each with a table and a concrete ask — and a “do not re-demand” list of checks already run.

Verify one slide against the paper before you trust the deck — same habit as Practice A.

■ If It Drifts, Steer

Faithful and fair are the two failure modes to police.

Symptom

- a confident statistic, no table
- a referee that only attacks
- a “major” with no clear ask

Say

- “Source every number or cut it.”
- “Add credit first, and be fair.”
- “State what evidence would satisfy you.”

■ Stretch, Then Hand Off

Push one artifact from “done” to “sharp.”

- Stretch: “Tighten Major 1 to three sentences and add one backup slide with the detail.”
- Stretch: ask the talk to add a single “limitations (author-flagged)” slide — honesty reads well.
- *Next — Practice C: make this paper compound, in an LLM Wiki.*