

CLASS 3 · PART A

Wiki LLM


Turn scattered papers into knowledge you can find and check.

Jiaqi Shao · NUS Business School



■ Build a Research Wiki That Remembers

No live demo. You will see the finished files and the commands that made them.



Then: what Sessions 1B, 2 and 3A add up to — and an appendix of five quizzes.

One idea per slide. If a slide needs two, it is two slides.

PART 1

The Problem

- *why every hand-kept research wiki dies.*



■ Remembering a Number Is Not Enough

One sentence you remember, and one question you cannot answer.

“Managers are **13% less likely** to answer AI-styled questions.”

Lin, Ni, Su & Xu (2026) — the paper you have had since Class 1B

Thirteen percent *of what?*

Take ten seconds. Almost nobody can answer this from memory.

■ Table 6 Tells Us What 13% Means

78 pages. 11 tables. 9 figures.

- | | |
|----------------------------|--------------------------|
| 1 Sample Screening | 7 Heterogeneity |
| 2 Platform Outline | 8 Sentiment Manipulation |
| 3 Design of Treatments | 9 Taken More Seriously? |
| 4 Summary Statistics | 10 Market Responses |
| 5 Balance Checks | 11 Crowding-Out Effects |
| 6 Likelihood of Responding | |

Which table shows whether this means 13 points — or 13% of 74%?

Table 6: $74\% \rightarrow 61\% = -13$ points; compared with 74%, the fall is about 18%.

The abstract gives the headline. Table 6 gives the exact meaning.

■ Four Things Go Wrong at Once

This is not a filing problem. Filing would not have helped.

The number is hard to check — which table supports it?


Materials scattered — PDF folder, highlights app, old chat, loose notes

Cross-paper ideas disappear — what did you learn across papers?

Conflicts disappear — does the next paper agree with this one, or not?

We need one place for sources, conclusions, links, and different readings.

■ Memex: Linked Knowledge Before Computers



Bush imagined a personal archive where readers could **save paths between related documents** and return to them later.

Memex gave us linked memory. It did not make the links easy to keep up to date.

■ The Idea Was Easy; The Upkeep Was Hard

Not the reading. Not the thinking. The bookkeeping.

“The tedious part of maintaining a knowledge base is not the reading or the thinking — it’s the **bookkeeping**. Humans abandon wikis because the maintenance burden grows faster than the value.”

— Andrej Karpathy, the LLM-Wiki gist

Every hand-kept wiki you have started died of exactly this.

The hard part was never the idea. It was keeping everything up to date.

■ LLMs Handle Upkeep, Not Judgement

It is good at precisely the half you are bad at.

“LLMs don’t get bored, don’t forget to update a cross-reference, and can touch 15 files in one pass.”

— Karpathy, same gist

Only you can

pick the sources · ask the question ·
rule on a disputed reading

Only it will

re-read everything monthly · keep 30
links consistent · never get bored

You make the decisions. The LLM keeps the routine going.

PART 2

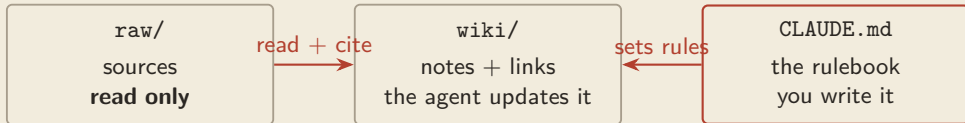
The Main Idea

- *three parts, one firm rule, three actions.*



■ Keep Sources, Notes, and Rules Separate

Keep what the source says separate from what you think it means.



Plus two bookkeeping files: `index.md` (contents) and `log.md` (timeline).

raw/ never changes · the agent updates wiki/ · you write the rules.

■ Why raw/ Is Never Edited

Follow one rule, and you can rebuild everything else.

If the agent may edit a source, then a source can no longer *settle* a dispute
— because you can never be sure it still says what the author wrote.

You can rebuild the notes. You cannot replace the original sources.

Delete wiki/, keep raw/ and CLAUDE.md, then rebuild in five minutes.

This is a safety rule built into the folder: the agent must never change the sources.

■ Three Actions: Ingest, Query, and Lint

Learn three verbs and you have learned it.

ingest — read a source, write a note with its source and page, update index and log

query — answer from the wiki and cite the notes; save useful answers

lint — list broken links, unlinked notes, missing sources, and unchecked pages

Ingest connects a new paper to what you already know.

■ A Source and Page Let You Check a Note

Five fields answer the question you could not answer ten minutes ago.

```
---  
type: literature-note  
status: needs-review  
source: raw/lin-2026-ai-ears.pdf  
pages: Table 6  
---
```

What the note should say — 13 percentage points; about 18% compared with the old 74% rate.

The shorthand is ordinary. The *note* that copied it without a page is what failed.

Frontmatter tells you where to look. You still decide whether the note is right.

SET-UP

Build It in Six Steps

- *everything you need is a folder and one markdown file.*



■ Step 0 — Make the Folder

The vault is just files. Obsidian is optional; the agent is not.

```
mkdir -p my-wiki/raw/clips my-wiki/wiki  
cd my-wiki && git init  
touch CLAUDE.md index.md log.md
```

Open the *same folder* in Obsidian (“Open folder as vault”) and in the terminal with `claude`.

Obsidian’s vault = Claude Code’s working directory. One folder, two windows.

■ Step 1 — Write the Rules (CLAUDE.md)


Copy this almost word for word. It is loaded every session.

```
# LLM Wiki --- rules
raw/  read-only sources; never edit, rename or delete
wiki/ maintained by you: literature, concepts, entities
Frontmatter on every page:
  type, status, source, pages, updated
Link recurring concepts with [[wikilinks]]
Update index.md and log.md after every ingest
```

One page of rules is what turns upkeep into a habit instead of a chore.

■ Step 2 — Put Sources in raw/

Use two or three papers you know well, so you can spot a wrong answer.



PDFs, web clips, notes — drop them in and stop. No renaming, no filing system.

The only habit that matters: a new source goes in raw/, never anywhere else.

■ Step 3 — The First Command

Plain English. There is no syntax to learn.

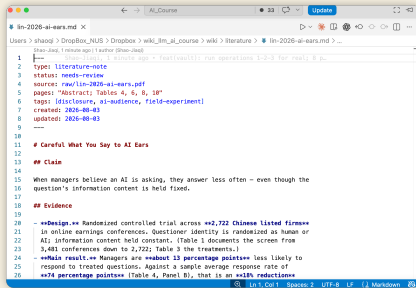
```
Read CLAUDE.md, ingest @raw/lin-2026-ai-ears.pdf,  
update index and log, then show the evidence path for each claim.
```

Four files must show the work — and one thing must not change:

a new page · an index line · a log line · **raw/ untouched**

The last clause of the command is the whole point: *show the evidence path.*


■ Step 4 — Check the Four Files



```
lin-2026-ai-ears.md X
Users > shaoqi > Dropbox_NUS > Dropbox > wiki_bin_course > wiki > literature > lin-2026-ai-ears.md > ...
1 1--- lin-2026-ai-ears.md - FastVault: run operations 1-3-3 for real; 8 p...
2 type: literature-note
3 status: needs-review
4 source: raw/lin-2026-ai-ears.pdf
5 pages: "Abstract; Tables 4, 6, 8, 10"
6 tags: [disclosure, ai-audience, field-experiment]
7 created: 2026-08-03
8 updated: 2026-08-03
9 ---
10
11 # Careful What You Say to AI Ears
12
13 ## Claim
14
15 When managers believe an AI is asking, they answer less often – even though the
16 question's information content is held fixed.
17
18 ## Evidence
19
20 -- **Design.** Randomized controlled trial across **2,722 Chinese listed firms**
21 in online earnings conferences. Questioner identity is randomized as human or
22 AI information content held constant. (Table 1 documents the screen from
23 3,481 conferences down to 2,722; Table 3 the treatments.)
24 -- **Main result.** Managers are **about 13 percentage points** less likely to
25 respond to treated questions. Against a sample average response rate of
26 **74 percentage points** (Table 4, Panel B), that is an **18% reduction**.
```

The work is done when you can open the files — not when the agent says it is done.

■ Step 5 — What You Now Have on Disk



Plain Markdown in a plain folder. Nothing here is locked into a product.

■ Step 6 — Lint, and Read the Warnings

An honest report flags problems. Warnings mean it is working.

```
$ python3 lint.py
Lint - 2026-08-05
OK    8 pages, 30 links
WARN  unreviewed: 5 pages
WARN  orphan: session-3a-takeaways
```

That orphan was real — one backlink fixed it; 29 links became 30.

A lint that always says “all clear” is the one to distrust.

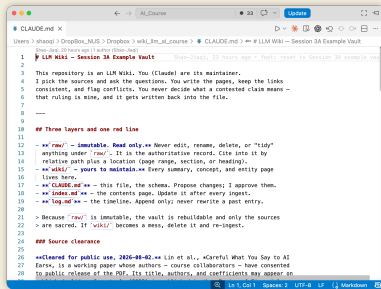
RESULTS

What It Produced

- *finished files, not a live demo.*



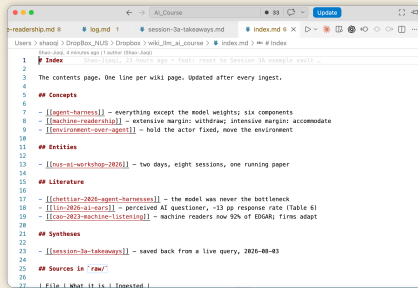
■ A Useful Note Always Points Back to Evidence



```
Users > shaoqi > Dropbox > wiki_rm_ai_course > CLAUDE.md > # LLM Wiki - Session 3A Example Vault
1 LLM Wiki - Session 3A Example Vault
2
3 This repository is an LLM Wiki. You (Claude) are its maintainer.
4 I pick the sources and ask the questions. You write the pages, keep the links
5 consistent, and flag conflicts. You never decide what a contested claim means -
6 that ruling is mine, and it gets written back into the file.
7
8 ---
9
10 ## Three layers and one red line
11
12 - raw/ - immutable. Read only. Never edit, rename, delete, or "tidy"
13   anything under raw/. It is the authoritative record. Cite into it by
14   relative path plus a location (page range, section, or heading).
15 - wiki/ - yours to maintain. Every summary, concept, and entity page
16   lives here.
17 - CLAUDE.md - this file, the schema. Propose changes; I approve them.
18 - index.md - the contents page. Update it after every ingest.
19 - log.md - the timeline. Append only; never rewrite a past entry.
20
21 > Because raw/ is immutable, the vault is rebuildable and only the sources
22 > are sacred. If wiki/ becomes a mess, delete it and re-ingest.
23
24 ### Source clearance
25
26 «Cleared for public use, 2026-08-02.» Lin et al., «Careful What You Say to AI
27 Firms», is a working paper whose authors - course collaborators - have consented
28 to public release of the PDF. Its title, authors, and coefficients may appear on
```

Every sentence on this page can be walked back to a location in `raw/`.

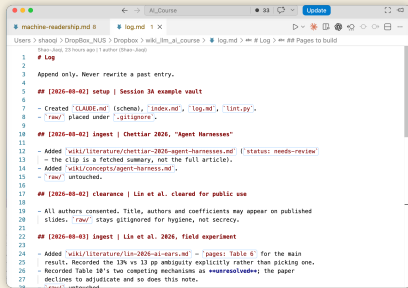
■ index.md — Find Anything by Reading



```
1 # Index
2
3 The contents page. One line per wiki page. Updated after every ingest.
4
5 ## Concepts
6
7 - [[agent-harness]] - everything except the model weights: six components
8 - [[machine-readership]] - extensive margin: withdraw; intensive margin: accommodate
9 - [[environment-over-agent]] - hold the actor fixed, move the environment
10
11 ## Entities
12
13 - [[nus-ai-workshop-2026]] - two days, eight sessions, one running paper
14
15 ## Literature
16
17 - [[chettiar-2026-agent-harnesses]] - the model was never the bottleneck
18 - [[lin-2026-ai-ears]] - perceived AI questioner, +33 pp response rate (Table 6)
19 - [[cao-2023-machine-listening]] - machine readers now 92% of EDGAR; firms adapt
20
21 ## Syntheses
22
23 - [[session-3a-takeaways]] - saved back from a live query, 2026-08-03
24
25 ## Sources in "raw"
26
27 | File | What it is | Indexed |
```

If you can read a shopping list, you can find your way around this file.

■ log.md — What Changed and When



```
1 # Log
2
3 Append only. Never rewrite a past entry.
4
5 ## [2026-08-02] setup | Session 3A example vault
6
7 - Created 'CLAUDE.md' (schema), 'Index.md', 'log.md', 'lint.py'.
8 - 'raw/' placed under '.gitignore'.
9
10 ## [2026-08-02] ingest | Chettiar 2026, "Agent Harnesses"
11
12 - Added 'wiki/literature/chettiar-2026-agent-harnesses.md' ('status: needs-review'
13   = the clip is a fetched summary, not the full article).
14 - Added 'wiki/concepts/agent-harness.md'.
15 - 'raw/' untouched.
16
17 ## [2026-08-02] clearance | Lin et al. cleared for public use
18
19 - All authors consented. Title, authors and coefficients may appear on published
20   slides. 'raw/' stays gitignored for hygiene, not secrecy.
21
22 ## [2026-08-03] ingest | Lin et al. 2026, field experiment
23
24 - Added 'wiki/literature/lin-2026-ai-ears.md' = 'pages: Table 6' for the main
25   result. Recorded the 13% vs 13 pp ambiguity explicitly rather than picking one.
26 - Recorded Table 18's two competing mechanisms as «unresolved»; the paper
27   declines to adjudicate and so does this note.
28 - 'raw/' untouched.
```

Six months from now this file tells you what you knew, and when.

■ The Wiki Can Connect Two Papers

Two sources went in. A third idea came out.

Cao et al. 2023 (RFS) — machine downloads 39%→92% of EDGAR; firms raise machine-readability ⇒ **accommodate**

↓
Lin et al. 2026 — perceived AI questioner, content held fixed; –13 pp response rate ⇒ **withdraw**

Fewer managers answer. Those who answer adjust how they respond.
In research terms: extensive margin = whether they answer; intensive margin = how they answer.

Nobody wrote that sentence into either paper.

■ Query — Asking the Wiki About This Course

```
# session-3a-takeaways.md X
Users > shaoji > Dropbox_NUS > Dropbox > wiki_fm_ai_course > wiki > syntheses > session-3a-takeaways.md > ...
Shao-Jiang, 6 minutes ago | 1 author (Shao-Jiang)
-----
1
2 type: synthesis
3 status: needs-review
4 sources: raw/clips/nus-ai-workshop-2026.md
5 pages: "session table; running-paper design" Shao-Jiang, 6 minutes ago • feat(wiki): run openai
6 tags: [teaching, takeaways]
7 created: 2026-08-03
8 updated: 2026-08-03
9
10
11 # Session 3A - Takeaways
12
13 *Saved back from a live query on 2026-08-03. This page is an answer that was
14 judged worth re-reading; two other queries that day were not saved.*
15
16 ## Claim
17
18 Three things survive the session, and each one has an artifact behind it that
19 the room watched being made.
20
21 ## Evidence
22
23 **1. 'pages:' is the difference between a citation and a memory.**
24 The abstract of \[\[lin-2026-ai-easy\]\] says "13% less likely"; Table 6 says 13
25 *percentage points*, which against a 74 pp base is an 18% reduction. One field
26 tells those apart and nothing else does - not memory, not a better search.
27 The source is not wrong; a note without "pages:" is.
28
29 **2. Compunding is countable.**
```

Every claim cites a page; where the wiki is silent, it says so.

■ Query — and the One Answer Worth Keeping

```
session-3a-takeaways.md X
Users > shaoqi > Dropbox_NUS > Dropbox > wiki_lm_ai_course > wiki > syntheses > session-3a-takeaways.md > ...

11 # Session 3A - Takeaways
21 ## Evidence
22
29 ==2. Compounding is countable==
30 The vault went from 2 pages / 2 links to 8 pages / ~38 links in one session.
31 Links grew faster than pages, and at least one of them =
32 [[teachise-readership]] - states a relationship that appears in xxxxxxxx
33 source paper.
34
35 ==3. Guardrails, verification, and you are three different things==
36 'raw/' being read-only is a guardrail. 'lint.py' is verification. Judging
37 whether a summary is «faithful» (w/ its source is neither - it is yours, and
38 'status: needs-review' is where that judgement is recorded as still owed.
39
40 ## Links
41
42 - [[lin-2020-ai-ears]] · [[car-2021-machine-listening]] - the two papers.
43 - [[teachise-readership]] - the cross-paper page that neither source contains.
44 - [[agent-harness]] - where guardrail vs. verification comes from.
45 - [[nus-ai-workshop-2026]] - where this session sits in the two days.
46
47 ## Open
48
49 - Point 2's link count is from a demo vault of four sources. Whether the growth
50 rate holds at forty sources is untested here.
51 - «What the wiki does not cover» nothing in this vault speaks to what
52 Sessions 3B, 4A or 4B teach. Those decks were never ingested.
53
```

Three queries asked. One written back. That ratio is about right.

■ More Links Make the Wiki More Useful

One session of work, measured.

Before

OK 2 pages, 2 links

After

OK 8 pages, 30 links

Pages $\times 4$. Links $\times 15$.

Links grow faster than pages — that is how the wiki becomes more useful.

The tenth paper you ingest is worth more than the first.

SESSIONS 1B · 2 · 3A

What They Add Up To

- *three messages, one arc.*



■ Three Messages, Three Sessions

The same three ideas, each one landing in its own session.

Agentic partner a junior collaborator that plans, acts, reflects — not a code generator

1B

Context engineering context quality caps output quality

2


Memory kept in files CLAUDE.md grows into a wiki that becomes more useful over time

3A

Today completes the third message.

■ What Each Session Handed You

Two sessions produce work. This one keeps it.



Sessions 3B, 4A and 4B all generate claims that need a source and a location.
This vault is where they should land.

If you keep one habit from the two days, make it: the source goes in raw/.

■ Three Rules Make the Wiki Reliable

You saw a file that proves each point.

1. pages: is the difference between a citation and a memory.

13 percent or 13 points — one field tells them apart, and nothing else does.

2. You can measure how the wiki grows.

2 pages / 2 links became 8 / 30. One of those links is in neither paper.


3. Safety rules, automatic checks, and human judgement do different jobs.

raw/ never changes. lint finds file problems. You decide whether the note matches the source.

Build the vault; let it grow with you.

■ A Five-Minute Routine Keeps the Wiki Alive

Five minutes a day, not a new job.



The practice kit `PC_practice.pdf` walks the whole loop again, slowly.

Ten papers in, it starts answering questions you did not ask it.

■ References

-  Vannevar Bush (1945). *As We May Think*. The Atlantic.
-  Andrej Karpathy (2026). *LLM Wiki* (public gist).
-  Lin, Ni, Su & Xu (2026). *Careful What You Say to AI Ears: A Field Experiment*.
-  Cao, Jiang, Yang & Zhang (2023). *How to Talk When a Machine Is Listening*. RFS.
-  G. Chettiar (2026). *Agent Harnesses: The Model Was Never the Bottleneck*.

Course materials: shao-jiaqi.github.io/ai-workshop-2026



Thank you

- *Build the vault; let it grow with you.*

APPENDIX

Five Quizzes

- *use with the practice session.*



■ Quiz 1 — Approve or Deny?

A real approval question while the agent checks the wiki.

Lint found: the abstract states “13% less likely”, but Table 6 reports 13 percentage points.

Proposed action: `edit raw/lin-2026-ai-ears.pdf` so the abstract matches Table 6.

Allow / Deny?

Vote — then name what tipped you off.

■ Quiz 2 — Which Lines Belong in CLAUDE.md?

The test: does every session need this line?

- a. "raw/ is read-only; never edit"
- b. "Today, focus on the Cao paper"
- c. "Frontmatter: type, status, source, pages, updated"
- d. "Use British spelling"

Pick the ones that earn a permanent line.

■ Quiz 3 — How Do You Record Disagreement?

Table 10 gives two possible reasons but does not choose between them.

A. Pick the more likely reason and write it as the finding

B. Leave the claim out of the wiki until someone settles it

C. Record both, mark the page status: `unresolved`, and say the paper declines to choose

Vote A, B, or C.

■ Quiz 4 — Lint Says “All Clear.” Is That Good?

Forty pages, one hundred links, zero warnings.

Your Friday lint returns OK 40 pages, 100 links and nothing else.

Is the wiki healthy?

Argue both sides in one sentence each.

■ Quiz 5 — What Is Wrong with This Note?

A page in your wiki says:

“Managers withdraw from AI audiences, so firms will disclose less as AI adoption rises.”

source: raw/lin-2026-ai-ears.pdf · pages: Abstract · status: verified

Name every problem you can see.

■ Practice Exercises

Do these with your own papers, in this order.

E1 — Build the folder and write `CLAUDE.md`. Have the agent read it back to you.

E2 — Ingest one paper. Check all four files, and checksum raw/ before and after.

E3 — Ingest a second paper. Find the first wikilink that crosses the two.

E4 — Write down what lint *should* find. Then run it. Compare.

E5 — Recover one claim from a wiki page back to the exact table in the PDF.

E5 is the final test. If you can do it, the system is working.