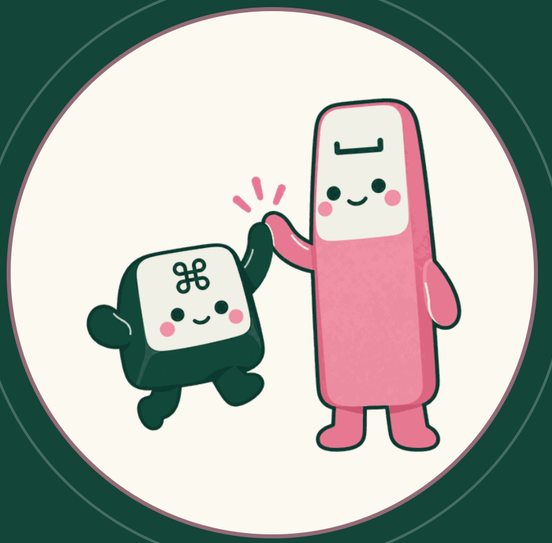
Practical AI at Work Same AI, Different Results

How finance professionals raise personal productivity with AI and agents



01

Delegate

a five-box work brief
instead of a question

02

Don't re-explain

keep your knowledge in files as context

03

Set standards

numbers with sources,
fixes become rules

Examples based on my own framework and public sources.

Last year, one demo got everyone nodding

This year, people use the same tools and get very different results.



01 **2025** After a demo, everyone thought 'that's how to use it'

02 **2026** Same tools, very different results

03 **Today's question** am I using it well enough not to fall behind?

KEY POINT

Keep one task of your own in mind as you listen.

Examples based on my own framework and public sources.

One fried chicken a month, do you spend it on AI?

A quick show of hands: what you spend on AI each month.

01	**₩0 a month**	I use a free account
02	**About ₩30,000**	I pay for one subscription
03	**Over ₩100,000**	I split work across several tools
04	**Your speaker**	over ₩2,000,000 a month

KEY POINT

Treat it as an investment in delegated work, not a cost.

Examples based on my own framework and public sources.

Three people using the same model

The gap opens in how you use AI, not in the tool.



01 **The asker** one question, one answer, starting over each time

02 **The delegator** hands over a whole task with the materials

03 **The standard-setter** passes on rules instead of fixing results

KEY POINT

Write down, in one word, which stage you are at today.

Examples based on my own framework and public sources.

Which box are you in right now?

Think about this week and pick the one box closest to you.

01	**Box one**	this week, I only asked AI questions
02	**Box two**	I attached materials and delegated a whole task
03	**Box three**	the second time, I didn't re-explain

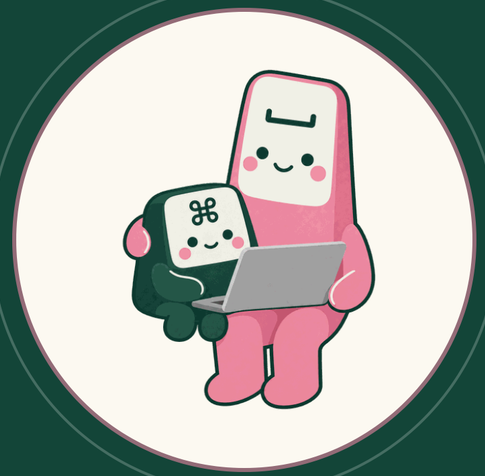
KEY POINT

Remember your box. I will ask again at the end.

Practical AI at Work

Difference 1
Delegate, don't ask

> Hand over work, not questions. We start with tasks from a finance professional's day.



A one-liner vs. a work brief

One-liner → work brief. The same task comes back differently depending on how you hand it over.

01	**"How's this company doing?"**	3 firms, 3 years of revenue and operating profit
02	**Generic answers every time**	source is DART annual reports
03	**No idea what it's based on**	one table and a 3-line memo on outliers
04	**Start over each time**	drop any number without a source

KEY POINT

Next time, state the shape of the output before you send.

ROLE BOUNDARY

People: goal and standards / **AI**: gathering and drafting

Examples based on my own framework and public sources.

A work brief needs just five boxes

The habit of filling the boxes beats prompt tricks.

STEP 01
Goal
what and why, who reads it

STEP 02
Materials
attach files and links

STEP 03
Output
fix the shape, like one table and a 3-line memo

STEP 04
Standards
numbers with sources, say so when unsure

STEP 05
Check
what you look at before approving

KEY POINT
Never leave the standards and check boxes empty.

Examples based on my own framework and public sources.

Work finance pros can delegate today

Meeting prep and report drafts work the same way.



01 **Filing comparison** several firms' financials in one table

02 **Research summary** reports and articles grouped by issue

03 **Email triage** long threads sorted into decided, open, owner

04 **Excel check** trace formula errors, rerun with new assumptions

KEY POINT

Start with work that takes half a day but needs only 10 minutes of judgment.

Examples based on my own framework and public sources.

Claude for M365 Claude for Outlook docs Claude for Excel docs

Fix the input contract before reading filings

We used figures presented in the 2025 filings of Samsung, SK hynix and LG Electronics.

01	**Scope**	FY2023–2025 consolidated; KRW million
02	**Reference**	6 metrics × 3 years × 3 firms; 54 source-checked values
03	**Preserve**	store source/reference inputs separately from model outputs

KEY POINT

Fix the scope first, then preserve inputs separately from outputs.

ROLE BOUNDARY

AI: process assigned sources / **People**: verify units, scope and originals

Assign different work to each model

These were different assignments, not a common-task model ranking.

01	**Evidence**	Jev selects / Gemini Flash extracts / Pro checks boundaries
02	**Analysis**	Astra interprets / Sol proposes and calculates scenarios
03	**Briefing**	Grok critiques / Opus drafts and rewrites after evidence restoration

KEY POINT

Specify the evidence passed to the next role, not just the role.

ROLE BOUNDARY

Models: distinct tasks / **People**: design handoffs and verification criteria

Examples based on my own framework and public sources.

Two models agree. Can we use it as is?

We checked extracted values, scenario amounts and evidence handoffs separately.

01	**First ask**	is model agreement the same as arithmetic verification?
02	**Extraction**	54 values matched; unsolicited industry hypothesis reviewed separately
03	**Arithmetic**	26/27 amounts matched; Grok missed a KRW 3 billion error
04	**Handoff**	narrow inputs mistaken for full source; evidence ledger restored

KEY POINT

Check agreement with arithmetic; check evidence-absence claims against originals.

ROLE BOUNDARY

Code: recalculate / **People**: check originals, preparation files and handoff scope

Inspect the evidence behind the answer

Recorded actual outputs and fixed-formula scenarios—not live inference.

01	**Records**	sources, selections, roles, verification and revised briefing
02	**Added**	what-if experiments and two dedicated 15-slide decks

KEY POINT

Open the inputs and corrections rather than copying a conclusion.

ROLE BOUNDARY

Web: recorded results and recalculation / **People**: verify and approve

Examples based on my own framework and public sources.

DART recorded runs · Dedicated poster · Dedicated brand

Today is about the people who use agents

We focus on handing work to agents and judging results, not on building them.

01	**Today**	handing your work to agents and judging the results
02	**Out of scope**	building the infrastructure and systems agents run on

KEY POINT

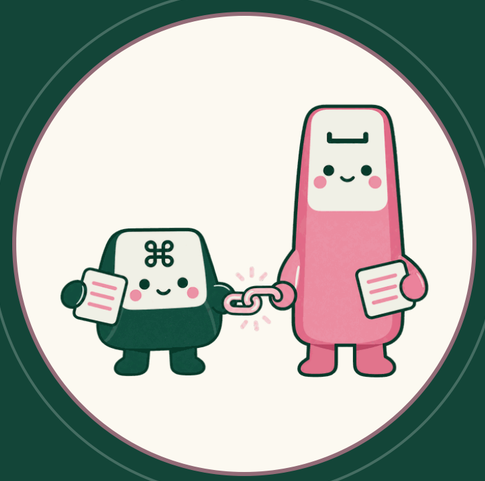
Start delegating what you can do from your own desk.

Examples based on my own framework and public sources.

Practical AI at Work

Difference 2
Stop re-explaining

> Productivity rests on knowledge management. When AI knows your context, you explain less.



The moment AI disappoints: 'Remember yesterday?'

Every time context breaks, the explaining starts over.



01 **Chat window** when the chat ends, so does the context

02 **App memory** convenient, but siloed in each app

03 **Files** you can attach them to any AI

KEY POINT

Put one thing you keep explaining into a file.

Examples based on my own framework and public sources.

Records to context in four steps

Delegated results flow back into your records, so the next explanation gets shorter.

STEP 01
Record
meetings, notes, judgments

STEP 02
Structure
Markdown plus metadata

STEP 03
Search
find only the pieces you need

STEP 04
Context
attach them and delegate

KEY POINT

Store delegated results back in your records too.

Examples based on my own framework and public sources.

An investor's personal knowledge base

Add fund notes plus terms and standards, and six boxes are enough to start.



01 **Manager notes** strategy, team, what you heard at past meetings

02 **Meeting notes** date, attendees, decisions, open items

03 **Judgment log** why you saw it that way, and whether you were right

04 **Work rules** briefs for AI and what it must not do

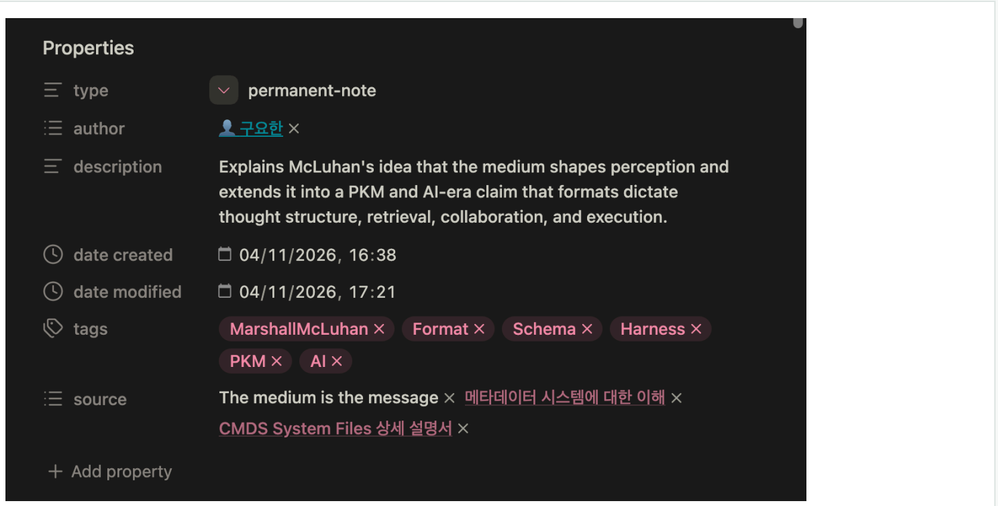
KEY POINT

One subject per note, and name the file after the subject.

Examples based on my own framework and public sources.

A few lines atop a note become search handles

A few properties tell AI which notes to pick.



Obsidian is a third-party app; this screen shows real CMDSPACE note properties. Captured 2026-09-10.

01 Type what kind of note it is

02 Description one line that tells AI whether to read it

03 Dates and tags when, and what it is about

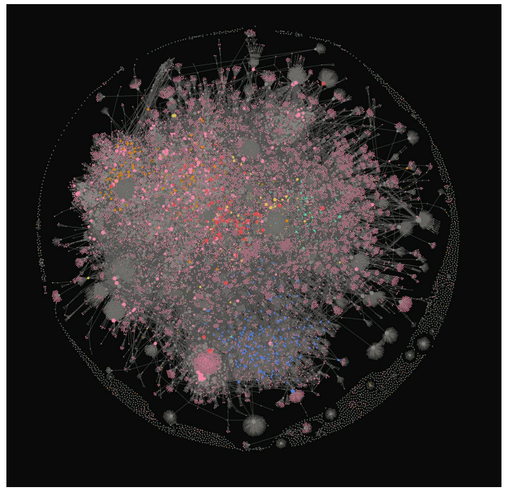
04 Source where this idea came from

KEY POINT

Add a one-line description to every note first.

Over four years of records are my AI's context

10,932 notes in the main vault, first entry May 2022 (measured Oct 4).



Obsidian is a third-party app; this is CMDSPACE's actual full graph. Note and tag names hidden. Captured 2026-09-10.

01 Dots one note each

02 Lines links between notes

03 Bright clusters topics often used together

KEY POINT

Start with one meeting note you write today.

Optional demo / may be replaced by this slide depending on the room.

First, separate company data from my knowledge

Today's scope is raising personal productivity within company policy.



01 **Company data** deals, commitments, client info only in approved tools

02 **My knowledge** public data, glossary, ways of working are mine to build

03 **Before attaching** is this file allowed to leave?

KEY POINT

Before attaching a file, ask once whether it may leave.

ROLE BOUNDARY

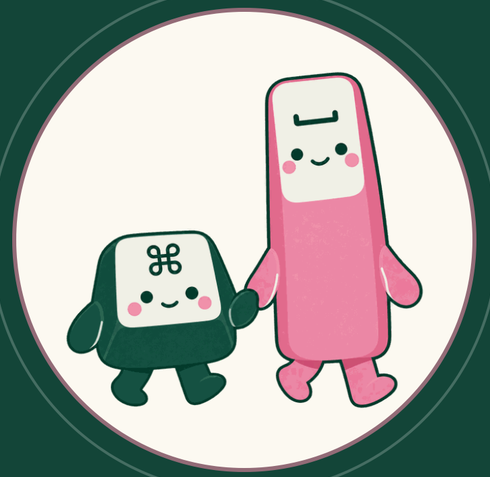
Company: deal and client data, approved tools / **Me**: public data and ways of working

Examples based on my own framework and public sources.

Practical AI at Work

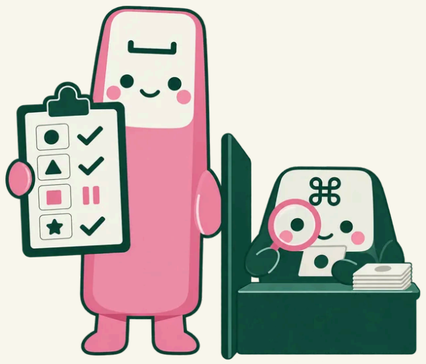
Difference 3
Standards, not reviews

> Rules for numbers and sources. Hand over the rules, not the fixes.



Reviewing is a bottleneck. Standards are not.

The faster AI gets, the more review piles up. A standard, once written, keeps working.



01 **Review** the faster AI gets, the more you have to read

02 **Standards** write them once and later results follow them

KEY POINT

Whatever you fix once, add as one line to your rules file.

Four standards for financial numbers

Numbers are where AI is wrong most convincingly.

STEP 01

Source

tag each number with its filing or file location

STEP 02

Units and periods

won vs. millions, consolidated vs. separate, fiscal year

STEP 03

Recalculate

totals and ratios by formula, not by AI

STEP 04

Unknown means blank

never fill in an unsourced number

KEY POINT

When numbers appear, ask for their source location first.

ROLE BOUNDARY

AI: find sources and build tables / **People**: totals by formula, check the blanks

Examples based on my own framework and public sources.

Claude for Excel docs

I didn't let it answer. I made it choose.

A real answer from my question desk, ask.cmdspace.work.



A real response to a real question, 2026-10-02 07:27 KST. Wording may differ each time.

01 A judge model picks the matching card from published service cards

02 Ask twice with the options reordered; trust it only when both agree

03 The answering model replies only from the chosen card

04 If it can't pick, it hands off to a person instead of making things up

KEY POINT

Don't ask AI to answer. Ask it to choose.

One page of standards changes the output

Without standards → With a standards file

01	**A different table layout each time**	Table layout stays fixed
02	**Number sources get mixed up**	Every number has its source location
03	**Units change sentence to sentence**	Units and period stated in one line
04	**My time goes into fixing**	I only make the call

KEY POINT

Attach the one-page standards file to every chat.

Examples based on my own framework and public sources.

Claude in Office October 2026

Available on paid plans (Pro, Max, Team, Enterprise); not on mobile apps.

01	**Excel**	generally available, cell cites and what-ifs
02	**PowerPoint**	generally available, slides from Excel data
03	**Word**	generally available, memo drafts from Excel
04	**Outlook**	beta, thread recaps and reply drafts
05	**Across apps**	turn it on and context carries over

KEY POINT

Start in the Excel and Outlook you open every day, not a new app.

From an Excel model to memo and slides

Requests adapted from the official docs' examples to investment work.

STEP 01

Read

How was this revenue calculated? It answers by pointing to cells

STEP 02

Change

Set the discount rate to 8%. Linked cells recalculate by formula

STEP 03

Carry over

Excel numbers into a draft investment memo and slide charts

STEP 04

Check

I match memo numbers to cell locations, then approve

KEY POINT

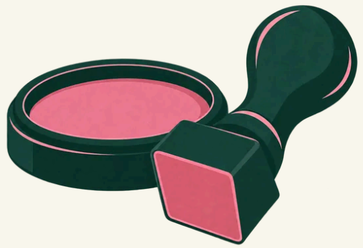
Get memo numbers together with their cell locations.

ROLE BOUNDARY

AI: explain calculations and draft / **People**: check cell locations and approve

What to check first at work

Add-in activity is not included in Enterprise audit logs.



01 **What you can try** practice on public data, you hit send, watch for hidden prompts

02 **What the company decides** deployment and approval, linking defaults, audit logs

03 **Client final versions** always reviewed by a person

KEY POINT

On a company account, check the admin settings first.

Even one vendor spans 100x

USD per 1M tokens, official prices as of 2026-10-06. The full table is at models.cmdspace.work.



01 **Fable 5.1** input 10, output 50, cache read 0.25

02 **Sonnet 5.5** input 2, output 10, cache read 0.10

03 **GPT-6 Astra** input 10, output 50, cache read 1.00

04 **GPT-6 Luna** input 0.1, output 0.5, cache read 0.01

KEY POINT

Give judgment to premium models and repetition to cheap ones.

Five habits to go deep while saving tokens

The app's effort dial is in the model picker.

STEP 01

One task, one chat

Long chats reread all earlier turns with every question

STEP 02

Match effort to the task

Low for tidying up, high for judgment and design

STEP 03

Judgment to premium models

Leave repetition and sorting to cheap models

STEP 04

Context in files

A repeated opening drops to the cache rate (API)

STEP 05

Split subscription and API

Subscription for work people do; API only for 24/7 jobs

KEY POINT

Keep one task per chat instead of one long chat.

Fireside chat

Your questions. Let's apply today's ideas to your own work.

It's not the model.
It's the structure.

Back to the opening question: which box are you in now?



01 **Delegate** a five-box work brief instead of a question

02 **Don't re-explain** keep your knowledge in files as context

03 **Set standards** numbers with sources, fixes become rules

KEY POINT

Pick just one of the three for this week.

Examples based on my own framework and public sources.

Three things to start tonight

You can measure where you stand at akm.cmdspace.work.



01 **Hand off one task** Delegate one recurring task this week with the five-box brief

02 **One note** Put what you explain every time for that task into one file

03 **Ask for sources** When numbers appear, add one line asking for their source location

KEY POINT

Tonight, write a brief for one recurring task this week.

Examples based on my own framework and public sources.

What you can take home

All free to use. Links are collected on the slides page.

01
AKM Index
self-assessment of AI knowledge management

02
Model guide
prices and uses by model

03
System Files
public version of my work rules for AI

04
LLM Wiki
an AI knowledge base built from your own materials

05
Book
Obsidian Professional Notes, Recommend, 2026

KEY POINT

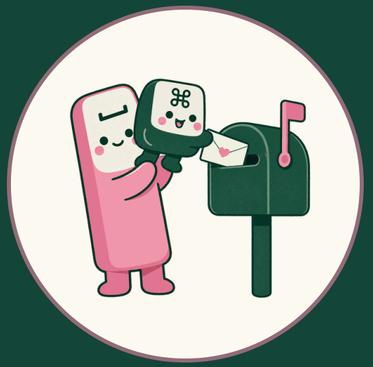
Send questions after the talk to ask.cmdspace.work.

Examples based on my own framework and public sources.

akm.cmdspace.work models.cmdspace.work system.cmdspace.work llm-wiki.cmdspace.work

Thank you

Materials and where to send questions.



Slides

labs.cmdspace.work/imm-1014/

Questions

ask.cmdspace.work

CMDSPACE

cmdspace.work

Yohan Koo

bio.cmdspace.work