



Homework Lab: Put the Four Big Moves Into Practice (90 minutes)

Why this matters

You'll leave the program with *concepts*. This lab turns them into **repeatable workflows** you can use tomorrow: verify outputs, guard against bias, train by doing, and work in a secure-by-default setup.

What you need to do this homework (pick one secure option)

- **Enterprise/privileged environment** (recommended): Microsoft Copilot (M365 E5/Purview), Lexis+ AI, Westlaw's AI tools, Casetext CoCounsel, or your firm's approved platform.
- **If you only have a consumer tool** (think ChatGPT or Gemini): Use **synthetic facts only** (no client identifiers). Turn off history/memory. Do not paste confidential information.

The Practice Scenario (Made-Up Facts for Hands-On Practice)

You're mediating a **business contract dispute** between Party A (supplier) and Party B (retailer) over late deliveries and withheld payments (\$180,000 at issue). Parties want speed, face concerns about cash flow, reputation, and ongoing business.

Exercise 1 — Make Verification Non-Negotiable (≈25 min)

Goal (the "why")

Experience the **full verify loop** so "no cite, no use" becomes muscle memory.

Task

1. Ask AI for a **neutral settlement options memo** (150–200 words) and **one sample clause** for a payment plan with milestones, interest, and non-disparagement.

2. **Verify** every factual/legal assertion and clause element.

Sample prompts (paste, then adapt)

Drafting prompt:

You are a neutral mediator drafting a brief options memo for both parties. In 150–200 words, list 3 settlement structures (lump sum, milestone plan, escrow). Then draft a one-paragraph payment-plan clause with milestones, dates, interest, and a mutual non-disparagement paragraph. Use plain language. Do not assume any admissions.

Self-check prompt (run after):

Identify every factual or legal claim in the memo and clause that requires verification. List them as checkboxes with the evidence or authority needed to confirm each one.

Verification checklist (save this)

- ☐ Dates, amounts, math (milestones total to \$180,000; interest math correct)
- ☐ Clause completeness (default/acceleration, notice method, cure period)
- ☐ Jurisdiction fit (non-disparagement/enforceability red flags)
- ☐ Citations confirmed in an **authoritative source** (KeyCite/Shepard's or firm forms)
- ☐ Final human read-through: clarity, neutrality, no hidden assumptions

Deliverable: A marked-up memo/clauses with track-changes and the completed checklist attached.

Exercise 2 — Install Bias Guardrails (≈20 min)

Goal (the “why”)

See how **demographics and phrasing** can skew advice—and how to neutralize it.

Task

Run the same option-generation twice: first with demographics, then anonymized. Compare.

Sample prompts

Version A (with demographics):

Party A is a woman-owned supplier; Party B is a national retailer. Suggest three settlement structures and propose initial monetary ranges for each option. Explain your rationale briefly.

Version B (anonymized/neutral):

Two commercial parties, no demographics provided. Suggest three settlement structures and propose initial monetary ranges for each option. Explain your rationale using only the facts provided (amount in dispute, delivery delay, ongoing relationship).

Bias scan (60-second check):

- ☐ Are ranges or rationales materially different A vs B without factual cause?
- ☐ Is tone/word choice more favorable to one side?
- ☐ Does advice rely on implied demographics?
- ☐ **Counterfactual test prompt:**

Re-evaluate your recommendations. Remove any demographic inferences. Base ranges solely on claim size (\$180,000), timing, and relationship goals. Explain changes.

Deliverable: A one-page comparison (what changed, why it matters, what you'll say in-session to keep it neutral).

Exercise 3 — Train by Doing, Not by Talking (Mini-Workflow) (≈25 min)

Goal (the “why”)

Practice the **end-to-end**: draft → verify → **redact** → **disclose**.

Task

Produce a 1-page **MOU summary** and label it correctly.

Sample prompts

First draft:

Draft a 1-page, neutral MOU capturing: (1) milestone payment plan totaling \$180,000; (2) a 30-day cure period; (3) mutual non-disparagement; (4) future purchase commitments subject to on-time delivery. Plain English, no legal conclusions.

Redaction helper:

Identify any details that could identify real parties. Replace with generic labels (Party A/B), and propose any additional redactions needed for privacy.

Disclosure language (paste into footer):

AI-Assisted Drafting Notice: Portions of this document were prepared with the assistance of an AI tool in a secure environment. All content has been reviewed, verified, and finalized by the mediator. No confidential client data was shared with non-enterprise systems.

Deliverables:

- The MOU (clean), plus your verification checklist and a **redacted** draft showing what you removed and why.

Exercise 4 — Go Secure-by-Default (Setup + Policy) (≈20 min)

Goal (the “why”)

Build a **repeatable safe setup** you can use on live matters.

Task

Complete this **security quick-start** and create a 1-page AI Use Policy for your practice.

Secure setup steps

- ☐ Use **enterprise** tools (M365/Copilot, Lexis+ AI, Westlaw, CoCounsel)
- ☐ Turn **memory/history OFF** (or use zero-retention mode) for prompts
- ☐ Apply a sensitivity label (“Confidential—Mediation”) and store in your **DLP-protected** workspace
- ☐ Anonymize prompts (Party A/B), remove PII and unique identifiers
- ☐ Keep an **audit log** (date, tool, prompt family, verification steps)

Policy template (paste and tailor)

Purpose. Enhance drafting/analysis while preserving neutrality, confidentiality, and accuracy.

Allowed uses. Brainstorming, plain-language summaries, clause variants, formatting, calculations.

Prohibited uses. Uploading client identifiers to consumer tools; relying on uncited/unchecked output; using AI to grade or rank parties.

Verification. “No cite, no use.” Every output checked against authoritative sources; discrepancies documented.

Bias guardrails. Anonymize prompts; run swap tests; perform a 60-second bias scan on key deliverables.

Disclosure. Use footer notice for AI-assisted drafts; client disclosures per jurisdictional rules.

Security. Enterprise platforms only; DLP, labels, and audit logs required.

Deliverable: One-page policy PDF and a screenshot of your security settings (or a written description if screenshots aren’t permitted).

Quick reference: Prompts you can reuse

- **Explain limits:**

Before answering, list 5 ways your output could be wrong or biased given my prompt. Then proceed with your best answer, and flag any places needing human verification.

- **Source-first drafting:**

Provide the answer only as bullet points with embedded citations (case, statute, or authoritative practice source). If you can’t cite, label the item “uncited—needs human source.”

- **Clause audit:**

List common failure points for payment-plan clauses (default, acceleration, interest usury, notice, dispute resolution). Compare them to the draft and mark any gaps.

- **Bias guard:**

Remove demographic inferences. Re-write recommendations using only the explicit facts and shared objectives. Explain the rationale in neutral terms.

What you will gain from these homework exercises:

Exercise 1 — Make Verification Non-Negotiable

Deliverable: Marked-up memo/clauses + verification checklist

What this builds:

- Habitual “no cite, no use” verification for any AI-assisted draft.
- Skill spotting unsupported claims, math errors, and clause gaps.
- A defensible audit trail you can show clients/courts/supervisors.
- Faster, cleaner revisions because issues are surfaced systematically.

Exercise 2 — Install Bias Guardrails

Deliverable: One-page bias comparison

What this builds:

- Muscle memory for anonymizing prompts and running quick swap tests.
- A simple 60-second bias scan you can apply to any output.
- Language discipline—neutral tone, fact-based rationales, balanced options.
- A repeatable step that supports mediator neutrality and party trust.

Exercise 3 — Train by Doing, Not by Talking

Deliverables: Final MOU + redacted draft + checklist + disclosure footer

What this builds:

- End-to-end workflow discipline: draft → verify → redact → disclose.
- Practical redaction skills that prevent leakage of identifiers/privileged facts.
- Clear, client-ready AI disclosure language you can reuse.
- Confidence producing a neutral, plain-language MOU under time pressure.

Exercise 4 — Go Secure-by-Default

Deliverable: One-page AI policy + security settings description

What this builds:

- A firm-ready AI use policy (allowed/prohibited uses, verification, disclosure).
- Set-up competence: zero-retention settings, labels/DLP, audit logging.
- Consistent team behavior and fewer ad-hoc, risky tool choices.

- Documentation that satisfies oversight, compliance, and client expectations.
-

Pro tips (time savers)

- Build a **Checklists** folder: Verification, Bias Scan, Disclosure Footer, Redaction Rules.
- Save your best prompts as “**Prompt Snippets**” you can paste into any tool.
- Track two metrics for 30 days:
 - **(a)** % of AI outputs needing corrections;
 - **(b)** issues caught by bias scan. Use this to show progress and coach your team.