



SharedIntel AI
Democratize Knowledge Together



Perplexity AI 101 - Q&A Session


Your Questions Answered

Live Q&A Follow-up from September 12th, 2025



September 12, 2025 | 11:00 AM EST

Hosted by: [Emma Farquharson, Ph.D.](#)

 **Your Questions, Expert Answers:** Based on the most common questions from our September 8th webinar attendees, this comprehensive Q&A provides practical, actionable guidance for implementing Perplexity AI safely and effectively in your business.

1

I'm starting on the free plan, so how can I do that safely—how do I protect

sensitive info?

What it means

You're dipping your toes into Perplexity's free tier—that's smart—it's powerful for quick research and doesn't cost anything. But because Perplexity does collect certain usage data (like prompts and uploads), it's wise to treat the free plan as a **public-facing research assistant**, not your secure document vault.

Why it matters for you


You're managing client data, internal priorities, and reputational risk—all at once. A single accidental upload or overshare can compromise confidentiality or trust. Starting safely lets you explore the tool's power *without exposing sensitive material*.

How to do it, step-by-step

- 1 Abstraction first.** Instead of copying "Client: Acme Corp Q2 budget review," say something generic like "Company A Q2 budget concerns." You keep the logic, drop the identity.
- 2 Avoid uploads at first.** On the free plan, anything you upload in a thread stays there for a while—unless you delete it. So work off URLs, public pages, or paste only sanitized snippets.
- 3 Use Incognito mode for sensitive areas.** It doesn't stop collection entirely, but it prevents your query history from being tied to your account—so you can test questions without leaving a trace.
- 4 Opt out of training.** In Settings or Preferences, there's a toggle to keep your prompts and outputs out of future training datasets. Flip it to "off" before your first question.

5 **Delete test threads.** When you're done experimenting, delete the entire thread—this ensures any pasted content or uploaded context is removed from your history.

6 **Web-only queries are your friend.** Default to using public URLs and filters (like Web, Academic, or Finance) so you always stay in safe territory until you're ready to level up.

 **Next action for you, right now:** *Open Perplexity's free plan and ask, "What is the typical budget margin for software implementation projects in 2025?" Use generic terms. Click the sources. No uploads, no risk, just a fast, verifiable insight.*

2

Perplexity vs. ChatGPT vs. Claude — which tool, and when?

What it means

You don't need more tools; you need the right one for the next deliverable.


Rule of thumb (remember this)

- **Perplexity** → "What happened recently? Show me sources." Use it for market scans, competitor moves, regulatory changes, and anything you must **forward with citations**.
- **ChatGPT** → "Help me create/brainstorm/iterate." Use it for creative drafting, campaign ideas, pitch outlines, and multi-step reasoning; **turn browsing on** when you need live sources.

- **Claude** → "Analyze this huge doc and write formally." Use it for long PDFs and measured, structured prose (policies, board memos). Then bounce back to **Perplexity** to ground your draft with **Academic/Finance** citations.

How to decide in 30 seconds

- 1 If your stakeholder will ask "Where did this come from?" → **Perplexity** (with the right *Focus* filter).
- 2 If tone, narrative, or brainstorming is the job → **ChatGPT** (browse only if needed).
- 3 If you're summarizing **long** or **technical** source material → **Claude**, then verify with Perplexity.

 **Try this:** Run your next brief in Perplexity (Finance), copy the outline into Claude for a polished draft, then paste back to Perplexity and ask: "Re-check every claim and attach numbered footnotes."

"Is this trustworthy enough to forward?" — your audit checklist

Why this matters

Two weak links and your deck loses credibility.


Four-part check (use this every time)


1 **Pedigree:** Do top citations include **.gov / .edu / regulator / IR pages / established publishers**? If not, re-run with **Academic** or **Finance**.

2 **Freshness:** Open at least **two links** and confirm dates align with the claim (pricing, law, security = strict recency).

3 **Agreement:** For any number that will travel, ensure **2+ independent sources** agree.

4 **Gaps:** If a sentence has no footnote—or the link doesn't support it—regenerate: *"Re-answer with **numbered footnotes aligned to claims**; prefer **Academic/Finance** sources from the last 24 months."*

 **Tip for executives/boards:** Lead with the claim, then show a short citation list (titles + dates). Perplexity's inline footnotes make this easy.

 **Try this:** *Take one slide you'll send today. Click two footnotes. Do the date and claim match? If not, re-run with Academic and fix it before you forward.*

4

"I uploaded a PDF—how do I remove it everywhere?"

Why this matters

Sensitive files shouldn't linger.

Clean-sweep sequence (do all four)

1 **Thread:** Delete the **message** that used the file—or delete the **entire thread** (safest in attachment-heavy chats).

2 **My Files:** If you stored it there, **manually delete** (persists until you do).

3 **Spaces:** **Manually delete** in every Space it touched.

4 **Pages/Links:** If you made a **Page** or shared a **public link**, **revoke/delete** and re-publish a clean version.

⚠ **Reality check:** Granular removal inside long threads can be fiddly. If in doubt, delete the thread and rebuild the final deliverable from a Page or PDF without attachments.

🚀 **Try this:** *Pick one sensitive file from last month and run the full sweep: Thread → My Files → Spaces → Links.*

5

RAG in plain English—why it changes your workflow

💡 Why this matters

You need answers you can defend, not prose that "sounds right."

🧠 What RAG means for you

Perplexity **looks it up first** (retrieves sources), then **writes** using what it found, with **citations**. That yields:

- **Fresher facts** for time-sensitive decisions.
- **Fewer hallucinations** because claims link back to real pages.
- **Defensible deliverables** with footnotes your stakeholders can click.
- **Blend public + internal** (Pro/Enterprise) so answers reflect your playbooks and the live web.



Try this: *"Using Academic only, summarize the last 12 months of research on SMB AI adoption barriers in 150–200 words with 5 citations (titles + dates)."*



Mac vs. Windows — what's actually different (and how to stay sane)



Why this matters

Mixed-OS teams need consistent behavior.



Stable path for parity

1


Use the **web app** as your standard. New features (Labs, Deep Research, fine-grained file controls) land there first and behave the same on Mac/Windows.


2

Treat **desktop apps** and **COMET** as optional add-ons you enable **after** your policies are in place. Some releases are **Mac-first** and staged.

3

Mobile (iOS/Android) is great for queries and quick uploads; do precise **deletion/export** in the browser.

 **Rollout rule:** Standardize onboarding and training on the web interface so everyone sees the same thing, then layer desktop/Comet selectively.

 **Try this:** Have one Windows and one Mac teammate complete the same web task (Filter → Academic; convert to Page; export PDF) and confirm identical outputs.

7

One feature to try today: Source Filters

Why this matters

It's the fastest visible quality lift for business research.

What to do

- 1 Draft in **Web** to frame the answer.
- 2 Re-run in **Academic** (evidence) or **Finance** (filings/IR docs) to produce the **forwardable** version—with stronger footnotes and less cleanup.

When to use which

- **Academic:** clinical/policy/technical claims.
- **Finance:** earnings, pricing, corporate facts that must match filings.

- **Social:** sentiment discovery—pair with Web/Academic to validate.



Try this: Re-run yesterday's summary with Focus: Finance and compare the footnotes. Are your citations board-ready now?

8

Filter strategy—pairing "recipes" you can trust



Why this matters

The **audience** dictates the proof you need.




Recipes

- **Client board decks: Academic + Finance** → peer-review + official numbers.
- **Competitive analysis: Finance + Social** → filings + market sentiment.
- **Market entry: Academic + Social + Web** → theory + voice-of-customer + breadth.
- **Crisis monitoring: Social + Web** → real-time chatter + authoritative coverage.



Principle: Pick the filter that matches the audience's trust threshold, then pair for balance.

 **Try this:** "For an executive brief, re-run this analysis with Academic + Finance; include 5 citations (titles + dates) and a 1-line source rationale for each."

9

Model choice—leave "Best" on, switch when style or task demands it

Why this matters


The same facts can land very differently depending on voice and structure.

Defaults and detours

- **Best:** Good balance of speed/quality for most research.
- **Claude:** Long-form, **structured**, measured tone (policies, HR/legal-ish).
- **GPT (o-series/Omni):** Brainstorming, creative variants, multi-step reasoning.
- **Sonar (Perplexity):** **Fast, retrieval-heavy** answers with clean citations.
- **Reasoning models (o1, DeepSeek-R1):** Math/code/explanatory "show your work"—then ground with a Pro Search run.

Workflow I use

Generate findings in **Perplexity**, draft the executive narrative in **Claude** (for structure), then back to **Perplexity (Academic/Finance)** to nail the footnotes.

 **Try this:** Answer a policy question in Best, then switch to Claude with the same prompt: "Re-write as a two-page policy brief with headings." Finish by

10

Cleanest handoff: Thread → Page → PDF (keep every citation)



Why this matters

You shouldn't lose footnotes when you polish.



Three-step flow

1

Do exploratory Q&A in a **Thread** using the right Filter(s).

2

Convert to Page to auto-structure; citations carry forward.

3

Export as PDF for a clean, fixed artifact—or share the **Page link** internally.



Security tip: If you attached sensitive files during research, don't share the thread—share the Page/PDF and keep the research thread private.



Try this: Take a working thread, Convert to Page, add a short intro, then Export PDF and verify the footnotes open correctly.

11

Data hygiene defaults for a 20-50 person team (light compliance)



Why this matters

Small teams can have enterprise-grade discipline without enterprise overhead.



Free/Pro (pilot) — set these norms

1

Publish a one-pager: **Allowed topics**, **Banned uploads**, **Sharing rules**.

2

Opt out of model training by default. Keep threads **private**. Avoid **My Files/Spaces** for business content.

3

Do **quarterly purges** of old threads.



Enterprise Pro (scale) — flip the controls on

1

Enforce "**no training on org data**".

2


Set **short retention** for thread attachments (e.g., 7 days).

3

Gate **uploads/downloads** and connectors via admin permissions; turn on **audit logs**.

4

Treat **Spaces** like repositories with **owners** and a **purge cadence**.

 **Try this:** Draft your one-pager today. If you have admins, set a 7-day attachment retention and assign owners for each Space.

12

Perplexity's COMET after the August 2025 disclosure — safe-use rules that keep you out of trouble...

Why this matters

Agentic browsers are powerful *and* new. Treat them like power tools.

Operating rules (until your security team signs off)

- 1 Assume web pages can contain **hidden instructions**. Avoid "Summarize this page" on unfamiliar sites.
- 2 Keep **email, banking, internal apps** in a separate, traditional browser profile.
- 3 **Disable auto-actions and saved credentials** in COMET; require explicit confirmation for any login/download.
- 4 **Pilot in low-risk contexts** first; log what you test and review the results.

13

A repeatable "RAG discipline" for regulated/evidence-sensitive work

💡 Why this matters

Compliance isn't a vibe—it's a trail of evidence.

📄 Two-stage pattern

- 1 **Retrieval spec first:** choose **Academic** (policy/clinical) or **Finance** (filings); set a **date window**; require **≥3 independent sources**.
- 2 **Answer second:** generate a short, plain-English summary with **numbered footnotes**.

📝 Templates you can paste

Evidence brief: "Using **Academic** sources since **2023**, summarize **X** in **200 words** with **5 citations** (title + link + date). Prefer reviews/guidelines; **exclude blogs/social**."

Regulatory check: "Using **Finance** focus, summarize the current requirements for **Y** with links to primary filings and regulator pages."


✅ Quality gates

- 1 Verify **≥50%** of footnotes.
- 2 Confirm **dates** and **authorities**.

3 **Cross-check** with Web for recency conflicts.

4 Capture **screenshots** of key sources for audit.

5 For high-stakes decisions, require **second-reviewer + SME sign-off**.

 **Try this:** *Run the evidence brief template and attach it to your policy wiki with the verification checklist completed.*

14

Perplexity Labs — high-value mini-apps (and the governance to keep them useful)

Why this matters


You can produce lightweight dashboards and tools without new software—if you set guardrails.

Great first builds (public data only)

- **Executive KPI Brief:** weekly market + competitor + regulatory snapshot with charts.
- **Pricing Monitor:** tracked public plan tiers across competitors.
- **Policy Checker:** summarize changes and output a red/amber/green action list.

Governance that keeps quality high

- 1 Name an **owner**; lock sources to **public/premium** providers.
- 2 Set a **refresh SLA** (e.g., weekly).
- 3 Spot-check **3–5 citations per release**.
- 4 Limit who can **export/share** externally.
- 5 Use **Max** subscriptions sparingly (unlimited Labs, early features) for 1–2 power users building org-wide assets.

 **Try this:** Draft a one-page "Labs Charter" (owner, sources, refresh cadence, QA checks) and pilot a Pricing Monitor with public pages only.

15 Team rollout — good → better → best (30–60 days)

Why this matters

Adoption sticks when you scale by behavior, not by licenses.

Good (Weeks 1–2)

- 1 Train Source Filters + the **trust audit** checklist.
- 2 Publish the policy one-pager.


- 3 Identify **2–3 champions**; keep research **Web-only** at first.

Better (Weeks 3–4)

- 1 Champions on **Pro**; introduce **Pages** and one simple **Labs** project (public data).
- 2 Stand up a **Spaces taxonomy** (Research, Client, Ops, L&D) with **owners**.
- 3 Add the **RAG discipline** checklist to your wiki.

Best (Days 30–60)

- 1 Expand Pro; evaluate **Enterprise Pro** for admin controls and retention.
- 2 Promote 1–2 creators to **Max** if they're building reusable assets.
- 3 Run a monthly **citation audit** (3–5 deliverables) and a quarterly **purge** of stale Spaces/threads.

 **Success looks like:** Your team produces a citable brief in under 15 minutes, tab-chaos drops, and nobody's sweating data exposure.

Custom GPT Trained on Your Data!

**Exclusive
15% Discount!**

\$339

Why You Need a Custom GPT:

- **Strategic Insights:** Access actionable insights tailored to your business for smarter decision-making and growth.
- **Consistent Brand Voice:** Seamless output aligned with your brand's tone, style, and guidelines.
- **Dynamic Adaptability:** A Custom GPT evolves with your business, learning continuously to meet your changing business needs.
- **Efficiency and Scalability:** Faster, more relevant engagements, saving you time and maximizing productivity.

Elevate Your Business with Precision, Security, and Expertise!

With **SharedIntel AI**, you get:

- **Quality and Precision:** We use the right data for accurate, high-quality outputs.
- **Security and Privacy:** Secure data handling to protect sensitive business information.
- **Expert Onboarding:** A 30-minute Zoom walkthrough to ensure smooth orientation and immediate value.
- **Optimized Prompts and Use Cases:** Accompanying document with impactful ChatGPT prompts and practical use cases, streamlining your workflows.

Contact ***Emma Farquharson, Ph.D.*** to discuss how AI can meet your business needs today.

Email: contact@SharedIntelAI.com

Book a preliminary call:

<https://calendly.com/sharedintelai/30min>



SharedIntel AI



