

AI Workflow Training Guide

Preview Edition

A practical resource focused on workflows, tools, process, and education

About this preview

This file contains the opening sections of the guide. To request the complete version, email info@investmentagent.org.

1. Purpose & Design Principles

This manual is written for serious practitioners: independent analysts, research teams, educators, operators, and small organisations that want a disciplined way to use AI. It assumes a working command of structured analysis, evidence handling, and written communication. Most examples use public-company research because the source material is rich and easy to verify, but the workflows are designed to teach process, not to confine the reader to a single role or institution.

The four objectives are straightforward:

- **Industrialise AI use across your process.** Move beyond ad-hoc experimentation toward structured, repeatable workflows you can rely on whether you work solo or on a team. You should be running the same earnings pipeline each quarter, not reinventing it every time.
- **Standardise workflows.** Whether the task is a post-earnings note, an initiation, or a weekly sector review, you should have a codified template and tool chain. Standardisation reduces variance in output quality and makes the process much easier to improve over time.
- Exploit each tool's capability bias—but route by deliverable, not by brand. No single AI model should own an entire workflow. The right approach is to route each sub-task according to what the deliverable actually needs: live external facts, grounded document analysis, structured drafting, spreadsheet analysis, challenge, or automation.
- Enforce verification and governance. AI output is a first draft, never a finished product. Every number, every factual claim, and every data point must be verifiable before it enters a model, note, deck, or decision journal. For solo practitioners this means self-imposed discipline; for teams it means clear review rules and documented accountability.

2. Capability Bias: Who Does What Best

Before diving into exercises, it is essential to understand what each tool is optimised for. The following breakdown reflects real-world performance as of early 2026 and should be revisited quarterly as capabilities evolve.

2.1 Perplexity

Perplexity is the primary research engine for anything that requires real-time information with verifiable citations. Its core strength is that every claim comes with a source link, making it the natural first stop for pre-market briefs, live earnings reaction work, news scans, and fast background research. The 'computer' is a game changing agentic environment (not covered here though).

2.2 Claude

Claude is a strong workhorse for structured analytical output. Its Excel modelling capability can approach sell/buy-side quality on clean cases: it can build linked three-statement models, DCFs with scenario analysis, and sensitivity tables that are genuinely useful when carefully checked. It is also very effective for long-form drafting, presentation creation, and turning rough research into readable prose.

2.3 ChatGPT

ChatGPT is one of the central orchestration, synthesis, and quality-control layer. Its role is not limited to ideation. It is strongest when the task requires turning messy inputs into a structured work product: scoping a research plan, extracting evidence from filings and transcripts, comparing documents, selecting the right valuation framework, auditing spreadsheets, drafting review-ready summaries, stress-testing conclusions, and preparing follow-up diligence questions.

Projects make ChatGPT the natural home for each live company, sector, or recurring review process because chats, uploaded files, and custom instructions remain together in one persistent workspace. GPTs allow you to package repeatable workflows—for example, earnings triage, decision memo QA, model review, management-meeting preparation, or risk-monitoring—with instructions, reference files, approved tools, and versioned logic.

ChatGPT is not the source of record for filings, market data, or consensus numbers. Its role is to structure the problem, normalise the evidence, surface contradictions, convert raw material into decision-ready output, and make the user faster and more rigorous without displacing human judgment.

2.4 Gemini

Gemini excels at large-context research. When you need a comprehensive sector report that synthesises dozens of sources, or a regulatory landscape mapping that covers multiple jurisdictions, Gemini's large context window and grounded search capability make it the strongest option. Its Chrome and Workspace integration allows seamless multi-tab research and direct interaction with Gmail, Google Docs, and Sheets. It is particularly well-suited for broad market mapping, policy analysis, and thematic research where breadth of coverage matters more than depth on a single name.

2.5 NotebookLM

NotebookLM is a source-grounded knowledge base. You upload your own PDFs, documents, and URLs, and the model's responses are strictly grounded in those sources—it will not hallucinate facts from outside your document set. This makes it especially valuable for earnings transcripts, company decks, annual reports, expert interviews, or collections of industry reports that you want to interrogate with confidence.

3. Core Principles for Serious Analytical Work

These seven principles apply to every interaction with every tool. They are non-negotiable and should be internalised by anyone using AI for serious analytical work, whether you are operating alone or as part of a team.

Principle 1: All numbers must be verifiable.

Any EPS figure, TAM estimate, growth rate, valuation multiple, or funding number that originates from AI must be cross-checked against primary sources—company filings, terminal data, or premium research databases—before it enters a model, memo, or deck. This is not optional. AI models hallucinate numbers with

confidence. The [VERIFY] tag convention used throughout this manual exists to make unverified figures immediately visible so they are never accidentally published.

Principle 2: Role + constraints in every serious prompt.

A well-structured prompt contains six elements, and omitting any of them degrades output quality materially. Always specify: the role (e.g., “senior TMT analyst covering listed software”), the goal (e.g., “post-earnings note for my review process”), the context, the constraints, the required output format, and the standard of evidence.

Principle 3: Workflows beat one-off prompts.

The real productivity gains come from structured multi-step workflows, not isolated prompts. Think in control points: one tool gathers or grounds the evidence, ChatGPT organises the evidence and identifies what matters, another tool may draft or model where appropriate, and the analyst verifies and finalises. A representative workflow is: Perplexity gathers current external facts → NotebookLM tests those facts against the relevant source documents → ChatGPT converts the material into an evidence map, challenge log, and PM-ready summary → the analyst verifies and finalises. The desk should design workflows around handoffs, auditability, and repeatability rather than letting each analyst improvise from scratch.

Principle 4: Standardisation beats brilliance.

Shared templates, Skills, Spaces, Projects, and prompt repositories reduce variance in quality across a process or team. When you use the same earnings review template each quarter, you know exactly what to expect and where to find each piece of information. Standardisation saves time, improves comparability, and makes mistakes easier to catch.

Principle 5: Human judgment stays on top.

AI handles structure, drafting, data gathering, and boilerplate. Humans make the assumptions, assess edge cases, weigh qualitative factors, and set priorities. This division of labour is the entire point: you free up cognitive bandwidth for the work that actually requires judgment by delegating the mechanical work to tools that do it faster and more consistently.

Principle 6: ChatGPT should be the first and last pass.

Use ChatGPT at the beginning of a workflow to define the question properly: what evidence is required, what output is needed, what the likely weak points are, and which valuation frameworks fit the case. Use it again at the end to challenge the draft, surface contradictions, identify unsupported claims, and prepare the questions the PM is most likely to ask.

Principle 7: Every serious deliverable needs an evidence map.

Before drafting a note, memo, or deck, convert the raw research into a simple evidence structure: claim, supporting evidence, source type, confidence level, verification status, and kill condition. ChatGPT is particularly useful for building this structure quickly, but the analyst remains responsible for deciding whether the evidence is actually sufficient.

4. Structure of the Training

The training is designed as a progressive curriculum with four tiers, each building on the skills developed in the previous one:

- **Tier 1 — Micro Skills (30–60 minutes each, single tool).** These exercises teach you how to use each tool individually for a specific, self-contained task. The goal is to build fluency and muscle memory with each platform before combining them.
- **Tier 2 — Modules (2–4 hours each, two to three tools).** These exercises chain multiple tools together into realistic workflows—an earnings pipeline, a decision memo workflow, a deck-building process. The goal is to learn how information flows between systems and where each handoff adds value.
- **Tier 3 — Full Cases (1–2 days, full tool stack).** These are end-to-end simulations of major deliverables such as a full company deep dive. Every tool is used in sequence, and the output is a complete, review-ready deliverable pack.
- **Tier 4 — Automation & Recurring Tasks (always-on).** These exercises set up the infrastructure for ongoing, automated workflows—daily briefs, weekly monitors, assumption trackers—that run without manual intervention and keep you informed around the clock.

Each exercise specifies its objective, the scenario context, the detailed task, the prompt language, and a thorough explanation of what constitutes excellent output, including common pitfalls and how to iterate toward a better result.

Request the complete guide

The full guide includes the remaining micro exercises, multi-tool modules, full-case workflows, and the automation sections. To receive it, contact **info@investmentagent.org**