Individual Submissions — Reflections (150–200 words each)

Reflection A — Prompt Engineer

In this project, I learned that the quality of AI help depends heavily on how I frame prompts. Broad "how do I import a CSV" prompts produced generic advice that sounded right but missed key specifics in our file. When I asked targeted prompts about missing tokens, mixed date formats, and percent/currency fields, the responses became much more actionable. However, even then, AI tended to assume common missing tokens (blank/NA/N/A) and did not naturally include "." and a single space as missing—both of which appear in our dataset. That taught me to treat AI as a starting point, not a checklist.

My contribution was to iterate prompts until we had guidance on (1) what assumptions import tools make and (2) how to verify them with evidence. I also focused on prompts that forced a "report failures" mindset, such as quantifying date parsing success rates. Going forward, I will use AI to generate candidate risks quickly, but I will always pair that with a plan to test those risks on the actual dataset.

Reflection B — Statistical Auditor

This activity made auditing feel like real data science rather than "trying to catch AI being wrong." Many AI suggestions were reasonable, but they were not automatically trustworthy for our file. My main contribution was to turn advice into verifiable claims and check them using the dataset. For example, instead of accepting "declare missing values," I looked for what missing values actually look like and found multiple representations in the same column (blank, NA, N/A, '.', and a single space). Similarly, rather than assuming date parsing is straightforward, I checked formats and confirmed that `host_since` and `last_scraped` mix ISO and US formats.

I also focused on silent-failure risks: percent and currency stored as text, and encoding issues in names like Zoë and María. These problems don't always throw errors; they show up later as incorrect summaries, unexpected NA creation, or broken joins. I learned that a strong audit produces evidence (counts, examples, parse-failure summaries), not opinions. That is what makes the auditor role meaningful even if AI is usually correct.

Reflection C — Synthesizer

My role was to convert our prompts, audit evidence, and AI guidance into a coherent explanation and presentation. AI produced many correct-sounding tips, but they were scattered; my job was to organize them around the central idea that importing is a decision step with hidden assumptions. I emphasized the assumptions we could demonstrate with evidence: multiple missing-value tokens, mixed date formats, percent/currency strings, and non-ASCII text. I also made sure we didn't simply paraphrase AI. Instead, I wrote a human explanation that connects each assumption to a real consequence (misleading categories, silent NA creation, biased sample size, distorted summaries, or corrupted text affecting joins).

Building the slides forced clarity: we had to separate "what AI suggested" from "where it fell short" and then show our corrected understanding with concrete evidence. That structure made our learning visible and would differentiate groups even if they used similar tools. I left the audience with a practical routine ("import report" mindset) because that is what I would want as a future student.