

AI in the Real World

Data, Power, & Society - Day 10

Making Decisions with GenAI

CMSE 492 (aka CMSE 101) • Fall 2026

Prof. Danny Caballero

Monday's Activity, Today

The power outage cancelled Monday's class before we could run the budget-cut activity in person.

- We're running it as today's session.
- Everything below was originally written for Monday — the instructions haven't changed, only the day.

The Average Is a Decision

Ask a generative AI system almost anything and you get the "consensus", i.e., the middle of everything that it was trained on.

It's usually reasonable, and often useful. But that can be a problem.

A wrong answer invites you to check the work while a reasonable answer does not.

- **Michigan's MiDAS** (2013–2015): a rule-based automated system flagged unemployment fraud with no human review. 93% of the 22,000 determinations examined involved no fraud. ~40,000 people accused; some lost homes or filed bankruptcy.
- **Australia's Robodebt**: took annual income, divided it evenly across the year, and treated that average as fortnightly earnings. Fine for steady work; wrong by construction for casual or seasonal work. Wrongly recovered A\$700M+ from ~400,000 people.

The average described nobody in particular, and the harm fell on the people it described least.

Today: Making Decisions from Generated Output

This week we add **Actions**, completing the DTPA framework.

You'll be handed a scenario with a budget you cannot meet and six programs — **two of which must be cut**.

Three phases:

1. **Human Discussion Cut (Baseline)** — no AI
2. **GenAI Baseline Cut** — hand the same table to at least two GenAI tools
3. **GenAI Context Probe** — ask what the table might be hiding

Please use the Google Doc posted in In-Class Materials on D2L for your work.

Phase 1 – Human Discussion Cut (Baseline)

No AI allowed yet.

- Discuss with your group using only the provided scenario table.
- Decide which two of the six programs to cut.
- Write a **3–5 sentence rationale** for your group's decision.
- Note what missing information would change your mind.

Keep this rationale. You'll compare it against the GenAI outputs in the next two phases.

Phase 2 – GenAI Baseline Cut

- Feed your scenario table to **at least two** different GenAI tools (e.g., ChatGPT, Gemini, Claude).
 - You can do this individually to compare with each other.
 - These systems tend to acclimate themselves to your prior requests, so you are likely to get different decisions.
- Use the same table, same framing, for each tool. Don't add hints.
 - If the tool asks for more information, push it to make a decision and provide a rationale.
- Paste each tool's **verbatim** choices and rationale into your group's tab in the [Class Decision Google Doc](#) (MSU login required).

Compare: Did the AI tools agree with each other? Did either agree with your group's Phase 1 decision?

Phase 3 – GenAI Context Probe

Run a follow-up prompt asking the AI to identify **critical hidden legal, financial, denominator, or systemic factors** the table might be ignoring.

- Paste the verbatim response into your group's tab in the Class Decision Doc.
- Does the AI tell you anything your group missed in Phase 1?
- Does it tell you anything *it* missed in its own Phase 2 answer?

This is the step MiDAS and Robodebt skipped: nobody asked what the number was hiding.

Whole Class Discussion

Comparing the Cuts

- Which programs did **humans** cut most often? Which did the **AI tools** cut?
- Did your two AI tools disagree with each other?
 - If the "consensus" changes depending on the tool, or on what you asked it earlier, what is it a consensus *of*?
- Did anyone change their mind after seeing the AI output?
 - What persuaded you: the **reasoning**, or the **confidence**?

The Reasonable Answer Problem

A wrong answer invites you to check the work while a reasonable answer does not.

- Did any AI rationale sound so sensible that you didn't question it?
- Compare your Phase 1 rationale to the AI's:
 - Which one names the **people** affected?
 - Which one talks mostly about **dollars and efficiency**?
- In Phase 3, the AI found factors it ignored in Phase 2.
 - If it "knew," why didn't it use them the first time?
 - Or did it raise the issue and I pushed you to ignore it?
 - **Who is responsible for asking the second question?**

Actions & Accountability (DTPA)

- A council adopts the AI's recommendation, and a program's clients are harmed. **Who is accountable?**
 - The tool? The analyst who pasted in the table? The council?
- MiDAS had no human review. Your process had humans in every phase.
 - Is "a human was in the loop" enough?
 - What would **meaningful** review have looked like here?
- Which part of this decision, if any, would you delegate to GenAI? **Where is the line?**

One Rule

If a government office uses GenAI to support budget decisions, what is the one rule it must follow?

- Write yours down; we'll share a few.
- Bring this into today's Exit Ticket and your Case Study 4 reflection.

Reminders

Exit Tickets, D2L, and the Course Website

- Don't forget to complete today's Exit Ticket by 11:59pm (on D2L)
- Case Study 4 (Actions & Critical Literacy) and Individual/Metacognitive Reflection due **Sunday, Sept 27th at 11:59pm**
- Spend time this week on Week 5's preparatory materials (Pick an area/report to summarize and ask DTPA style questions)
 - Forum Post due 11:59pm **Sunday**

Your instructors have given detailed feedback on Week 3's assignments.