

AI ETHICS COUNCIL

CASEBOOK

Ten days of evidence, stakeholder thinking, debate, and a code people can actually use.

Council member information

Name	Class / period	Council team

Essential question: When an AI system can help some people and harm others, how should we decide what is responsible?

Your council job

Do not begin with “AI is good” or “AI is bad.” Trace the decision, examine who has power, verify the evidence, and write a rule that another person could actually follow and check.

INVESTIGATE	LISTEN	DECIDE
Find traceable evidence.	Represent more than one stakeholder.	Name the action, limit, and check.

Privacy boundary: Analyze systems and supplied cases. You do not need to share personal feeds, messages, faces, location history, or family experiences.

DAY 1 | CASE INTAKE: WHAT HAPPENED?

Document one AI ethics problem before deciding what should be done.

Case snapshot

AI system / decision	What went wrong	Who was affected	Who built, chose, or used it

Two-source evidence log

Source title + organization	Link / date	Fact the source supports	Evidence or opinion?

Source test: Can another person find this source? Does it identify who collected the information or made the claim?

Initial claim

AI ethics matters in this case because _____ was affected when _____.

DAY 2 | RECOMMENDATION SYSTEMS

Follow what the system is designed to maximize, then map who gains and who becomes less visible.

System goal

This system appears to maximize _____. I think that because _____.

Stakeholder map

Stakeholder	Direct or indirect?	Possible benefit	Possible harm	Power to change the system

Redesign ♦ tradeoff

Change one rule. Who gains? Who may lose attention, money, access, or voice?

DAY 3 | PRIVACY + SURVEILLANCE

Judge a system using data, consent, benefit, risk, and the ability to opt out.

What does the system collect?

Data collected	Who can access it?	Was consent clear?	Can a person opt out?

Two-sided analysis

Possible benefit + who receives it	Possible harm + who carries the risk

Safeguard

Keep one useful function, but add a limit, permission, deletion rule, appeal, or human review.

DAY 4 | MISINFORMATION ♦ LATERAL READING

Do not judge a claim by how polished it looks. Leave the page and verify it.

Evidence sort

Example	Real / generated / altered / uncertain	First clue	What still needs checking

Lateral-reading trace

1 · Leave the page	2 · Check the source	3 · Find independent coverage	4 · Trace media to the original

Verification decision

I would / would not share this claim yet because _____. The strongest evidence I found was _____.

DAY 5 | STAKEHOLDER HEARING

Represent the people affected by one AI decision, not only the person buying the system.

Scenario

What decision does the AI system influence? Who controls the data, rule, or outcome?

Six-stakeholder hearing record

Stakeholder	What they need	Benefit or risk	Power / ability to appeal

Council recommendation

We recommend _____ because _____. To protect _____, the system must _____.

DAY 6 | WORK IN THE AGE OF AI

Break a job into tasks before claiming that AI will “replace the job.”

Job ♦ task breakdown

Task	AI / human / both	Why?	Risk if assigned badly	Safeguard or skill

Future role

People in this job may spend less time on _____ and more time on _____. A skill that becomes more important is _____.

Avoid the shortcut: AI may change tasks, staffing, access, and working conditions even when the job title remains.

DAY 7 | DEBATE PREPARATION

Prepare the assigned stakeholder position. Your role is not a statement of your personal belief.

Assigned role + position

Stakeholder role	Position: allow / limit / pause / prohibit	What this stakeholder values most

Evidence + reasoning

Claim	Evidence / source	How the evidence supports the claim	Tradeoff you will admit

Counterargument

Some people argue _____. From this stakeholder's perspective, _____ because _____.

Opening statement

Write a 45- to 60-second opening with position, reason, evidence, and requested action.

DAY 8 | COUNCIL DEBATE

Listen for evidence, stakeholder impact, tradeoffs, and a workable action, not the loudest speaker.

Live evidence notes

Speaker / role	Strongest evidence	Stakeholder impact named	Question or response

Class decision

Option	My vote	Best reason	Biggest unresolved risk

Individual reflection

The argument that changed or strengthened my thinking was _____ because _____. The council should _____.

DAY 9 | WRITE A TESTABLE CODE OF ETHICS

A principle is useful only if it names who acts, what they do, and how someone can check it.

Five principles

Area	Who must act?	Required action or limit	How will people check it?	Evidence from the unit
Bias and fairness				
Privacy				
Misinformation				
Transparency				
Accountability				
Weak → testable: Weak: “AI should be fair.” Testable: “Before launch, the team must compare error rates across affected groups and publish the results.”				

Council title ♦ one-sentence purpose

DAY 10 | GALLERY WALK ♦ FINAL SYNTHESIS

Test whether another group’s principles are clear, evidence-based, and usable.

Peer feedback

Council / principle	A strength	A question or missing stakeholder	Suggested revision

Defend one principle

Our most important principle is _____ because the evidence from _____ shows _____.

What changed?

At the start I thought _____. Now I think _____ because _____. One question I still have is _____.

Final check: Your answer may be nuanced. The goal is not agreement; it is a decision that can be explained, challenged, and revised with evidence.

TOOLKIT | WORDS, STEMS + EVIDENCE CHECKS

Use this page when the issue feels complicated or the discussion gets stuck.

Core vocabulary

Word	Working meaning	Use it in a sentence
ethics	rules and reasoning about what people should do	
stakeholder	a person or group affected by a decision	
tradeoff	a gain that comes with a cost or risk	
consent	clear permission given with a real choice	
transparency	enough information to understand or question a system	
accountability	who must answer for the outcome and repair harm	
safeguard	a rule, limit, review, or appeal that reduces risk	

Complete sentence stems

- The system is designed to _____, which benefits _____ but may harm _____.
- From the perspective of _____, the most important need is _____.
- The evidence from _____ supports the claim that _____.
- A reasonable safeguard is _____ because it keeps _____ while reducing _____.
- I disagree with _____; however, the strongest part of that argument is _____.

Evidence ladder: Weak: “I saw a post.” Better: a traceable source. Stronger: multiple independent sources with methods, dates, and limitations.