Session 1 Handout: Fair and Accountable AI: Ethics and Bias

FAST-17 Responsible AI for All Staff · 7formations · Instructor: Wadood Chaudhary · All examples fictional · Not legal advice · Suggested reading: Svekis, AI Literacy for All Employees (2026), chapters 8 and 9

Why this session matters

Every employee who uses AI makes small ethical decisions: what to trust, what to paste, whether a result treats people fairly, and what to do when something looks wrong. You do not need to be a manager or a lawyer to make those calls well. You need a few principles, a quick check, and to know who to tell.

Key concepts

1. Four principles of responsible AI

Principle	In everyday terms
Fairness	People are not treated worse because of who they are
Accountability	A person or team owns the outcome, not the tool
Transparency	People know when AI affects them and can question it
Privacy	Personal information is protected and used only as needed

2. The five-step ethical check

1. **Stakeholders:** who is affected?
2. **Risks:** what could go wrong, and for whom?
3. **Transparency:** do they know AI is involved? Can they question it?
4. **Oversight:** who reviews, and who can override?
5. **Compliance:** which law or policy applies?

3. How bias gets in

Type	What happens	Example (fictional)
Historical	Data reflects past unequal treatment	Résumé model trained on mostly male past hires
Representation	Some groups are missing or rare in the data	Voice tool fails on strong regional accents
Measurement	The thing measured is a poor stand-in	Call quality judged by call length
Evaluation	Tested on the wrong people or conditions	Safety detector tested only on the day shift

Removing a protected attribute does not always remove bias: other fields (zip code, availability flags, career gaps) can stand in for it.

4. The law in Massachusetts (not legal advice)

- **G.L. c. 151B** prohibits discrimination in employment. A tool does not change who is responsible.
- The **Massachusetts Attorney General's advisory (April 2024)** says existing consumer protection, anti-discrimination, and data security laws apply to AI.
- Elsewhere, **NYC Local Law 144** requires bias audits for automated hiring tools, a sign of where rules are going.

Worked example

North Shore Builders' (fictional) AI scheduler never offers weekend overtime to workers with a childcare flag. The flag acts as a stand-in for family status. The right response: write down examples with dates, report through the named channel, and do not quietly change the data yourself. The company, not the software, is accountable.

Try this before session 2

- Run the five-step check on one AI use you have seen at work.
- Read Svekis chapters 8 and 9.