

# Responsible AI Decision Lens

*A framework by Soulful AI*

AI is powerful -- and easy to misuse without realizing it. The Responsible AI Decision Lens helps teams pause and ask the right questions before building or shipping.

*It's not about slowing down. It's about seeing clearly.*

## Use this lens to:

- Spot hidden risks early.
- Make human-centered trade-offs visible.
- Strengthen trust with users, partners, and regulators.

## Human Impact

**Core Question:** Who benefits, who's burdened, and who's left out?

**What to Look For:** Direct & indirect effects on users, employees, and communities.

**Guiding Prompts:** Who gains or loses agency? • Would I want this used on my family?

**Watch Out For:** Removing user choice • Manipulative interactions

## Fairness & Bias

**Core Question:** Could this create or amplify bias?

**What to Look For:** Data sources, labeling, evaluation sets.

**Guiding Prompts:** Who's missing from our data? • Does performance vary by group?

**Watch Out For:** Homogeneous data • No bias testing

## Transparency & Explainability

**Core Question:** Can people understand how it works and why?

**What to Look For:** Model interpretability, documentation, user messaging.

**Guiding Prompts:** Could we explain this to a regulator?

**Watch Out For:** Black-box decisions in high-impact contexts

## Privacy & Data Stewardship

**Core Question:** Are we collecting only what we need, with consent?

**What to Look For:** Collection, storage, deletion policies.

**Guiding Prompts:** Would users be surprised by how we use their data?

**Watch Out For:** Hidden secondary uses • Indefinite retention

## Accountability & Oversight

**Core Question:** Who's responsible if something goes wrong?

**What to Look For:** Ownership, escalation, review cadence.

**Guiding Prompts:** Who signs off before launch?

**Watch Out For:** No accountable owner • Ad-hoc governance

## Long-Term Consequences

**Core Question:** What happens if this scales or is repurposed?

**What to Look For:** Systemic, environmental, or societal effects.

**Guiding Prompts:** Could this be misused at scale?

**Watch Out For:** Dual-use potential • Large-scale disruption

---

## How to Use It

**Kickoff:** Add a 10-minute "Lens Check" to project starts or PRD reviews.

**Decision Docs:** Capture takeaways (e.g., "Medium fairness risk mitigated via data audit").

**Retros:** Revisit after launch to track real-world effects.

**Optional:** Score each lens (1-5) -> visualize risk levels with a simple color bar.

## Keep It Visible

- Pin it in your Notion sprint template.
- Print it beside your roadmap.
- Revisit it when uncertainty spikes -- not just at the end.

---

**Take the AI Readiness Quiz ->**

Created by Soulful AI -- soulfulcraft.ai