



MODEL DIGNITY CHECK

From Chapter 9 of *AI and the Art of Being Human*

WHAT IT IS

A pre-launch checklist for ensuring AI systems preserve human dignity and catch blind spots before they scale.

WHERE YOU WOULD USE IT

- Before any AI system or feature goes live
- When reviewing existing AI implementations
- During design phases, to catch problems early
- When updating or retraining AI models

THE TOOL:

Five Questions To Ask Before Launch

- **Who becomes invisible when we optimize?** Name specific people, not categories—“elderly residents in walk-ups” not “some users”
- **What “normal” is baked into the training data?** Every dataset tells a story about who matters—who’s reality shaped this system?
- **How does this perform for our most vulnerable users?** Test on edge cases—the users with least power, resources, or technical literacy
- **Can affected humans understand and contest decisions?** Opacity breeds distrust—is there a clear path to challenge the algorithm?
- **Does this strengthen or erode human agency?** Are we augmenting human judgment or replacing it?

HOW TO USE IT

Before launch, document written answers to all five questions. Be specific—vague answers hide real problems. If any answer troubles you, redesign before deploying. This isn’t a compliance checkbox but a discipline for catching what pure optimization misses. Run this check again whenever you update or retrain the system.