

Karim's AI Playbook — One-Page Cheat Sheet

Pillar	3-Second Litmus Test	Key Moves	Watch-outs
1. Problem Framing	"If it works, can we act on it?"	<ul style="list-style-type: none"> Anchor to one KPI (\uparrow revenue, \downarrow cost, \downarrow risk) Draft a <i>success counter-factual</i> ("What happens if we change nothing?") 	Shiny-object syndrome; projects with no business sponsor
2. Data Discipline	"Can we source, store and share it?"	<ul style="list-style-type: none"> Catalogue data lineage & quality Define lean <i>feature contracts</i> (type, range, refresh cadence) Automate PII redaction at ingestion 	Hidden licence terms; drift when upstream systems change
3. Model Selection	"Do we really need a rocket when a bicycle will do?"	<ul style="list-style-type: none"> Start with baseline heuristics \rightarrow classical ML \rightarrow LLMs only if justified Compare open-source to hosted APIs (cost, IP, latency) 	Over-fitting to historical bias; vendor lock-in on proprietary stacks
4. Evaluation & Explainability	"Would a non-expert accept the answer?"	<ul style="list-style-type: none"> Choose <i>business-aware</i> metrics (e.g. Δ CLV, precision@top-k) Use SHAP or feature-importance plots in stakeholder packs Build an <i>A/B kill-switch</i> from day one 	Metrics that optimise vanity (e.g. accuracy minus cost curve)

Pillar	3-Second Litmus Test	Key Moves	Watch-outs
5. Responsible & Ethical AI	“Could this harm a vulnerable group—or our brand?”	<ul style="list-style-type: none"> • Bias audits on protected attributes • Publish model cards & data sheets • Map to ISO 42001 / EU AI Act risk tiers 	Ignoring edge-cases that end up on the front page
6. Deployment & MLOps	“Can we roll back in under five minutes?”	<ul style="list-style-type: none"> • CI/CD with canary releases • Feature store + model registry • Near-real-time drift alerts (data and concept) 	One-off notebooks smuggled into prod; lack of version pinning
7. Governance & ROI Tracking	“Are value and risk fully visible to the board?”	<ul style="list-style-type: none"> • Quarterly scorecard: value delivered vs forecast • Cost-per-prediction & carbon-footprint dashboards • Independent model-risk review committee 	Success metrics defined post-launch; no budget for post-mortems