

COURSE 05 · BUSINESS TRACK

AI Strategy and Economics

What changes when knowledge work gets cheap, and how to plan for it

2 DAYS

14 CONTACT HOURS

HANDS-ON

A course for decision-makers on reading the evidence about AI capability, estimating its effect on costs, roles and business models, and turning that into a sequenced, measurable adoption plan with governance attached.

AUDIENCE

Founders, executives, managers, and product, strategy and operations leads.

DELIVERY

Participants work on their own organisation throughout: task inventory, exposure map, opportunity portfolio and roadmap.

CAPSTONE

A 90-day AI roadmap with a business case for each initiative, presented for critique.

Feasibility verdict: Feasible as a decision-making course when participants work on their own organisation and technical topics stay at the depth needed to judge them. Building is limited to demonstrations; teams that need to build should continue with Course 4 or Course 3.

Version 1.0 · October 2026 · dinu.at/teaching/ai-strategy

Course overview

How to work out which parts of your organisation AI will change first, where the value goes when that work gets cheap, and what to do about it in the next 90 days. Participants work on their own organisation throughout: task inventory, exposure map, opportunity portfolio and roadmap.

Learning outcomes

- Read capability evidence critically: what benchmarks such as GDPval and METR's task time horizons measure, what they leave out, and how a demo differs from a deployed system.
- Estimate how exposed the tasks in a role or function are, using frequency, digitisation, checkability and the cost of errors rather than job titles.
- Explain where value moves when knowledge work gets cheap: accountability, proprietary data, distribution, regulation, physical operations and customer relationships.
- Decide between buying, building and rebuilding software once agents make custom software much cheaper to produce.
- Assess the effect on pricing and business models, including the shift from seat-based to usage- and outcome-based pricing.
- Build a sequenced portfolio of AI initiatives with owners, costs, measurable outcomes and stop criteria.
- Plan upskilling and operating-model changes so that adoption doesn't depend on a few enthusiasts.
- Map initiatives to the EU AI Act, ISO/IEC 42001 and the NIST AI RMF, and name the risks that need a decision at board level.

Realistic depth within two days

IMPLEMENT	PRACTISE	SURVEY / DEMO
• Task inventory and exposure map of the participant's own organisation • A business case per initiative • A 90-day roadmap	• Reading benchmark and vendor claims • Build, buy or rebuild decisions • Pricing and business-model scenarios	• Model internals and training • Legal compliance work for a specific system • Hands-on agent and automation building

Conditions that make the schedule feasible

- Participants bring an outline of their organisation's main functions and processes; a prepared case company is available for anyone who cannot share internal information.
- Technical content is demonstrated live rather than built by participants.
- Legal topics are covered at the level of what a leader has to decide and ask; they don't replace legal advice.
- Roughly 6 hours of structured exercises on participants' own organisations, 6 hours of concepts and demonstrations, and 2 hours of presentations and critique.

Recommended participant profile

Primary audience. Founders, executives, managing directors, heads of function, and product, strategy, finance and operations leads who decide where AI is used and how it is paid for.

Prerequisites. No technical background is required. Participants should have used at least one AI assistant for their own work and know their organisation's main processes and cost structure.

Class size. 8–16 participants per instructor; in-house runs work best with a leadership team of up to 12.

Day 1: Capability, exposure and where value moves

Build an accurate picture of what AI can do now and what that means for this organisation.

Time	Module	Content and activity
09:00–09:30	Orientation	What has changed in the last two years, the course process, and the question each participant came with.
09:30–10:30	Reading the evidence	Benchmarks (GDPval, METR time horizons, coding and agent benchmarks), their blind spots and contamination, demos versus deployed systems, and how to question a vendor's claim.
10:30–10:45	Break	Coffee and questions.
10:45–12:15	Demonstration and review	An agent completes a realistic piece of knowledge work end to end: research, analysis, a document and a slide deck. The group reviews what would and wouldn't survive contact with a customer.
12:15–13:15	Lunch	Lunch break.
13:15–14:00	Which work changes first	Tasks, not jobs: frequency, digitisation, checkability, cost of errors and judgement. Why verifiable digital work moves first, and physical and relationship work later.
14:00–15:00	Exercise 1: exposure map LAB	Participants list the main tasks in their own organisation and score how exposed each one is.
15:00–15:15	Break	Coffee and questions.
15:15–16:15	Where value moves	When producing software and documents gets cheap, value moves to accountability, proprietary data, distribution, regulation, physical operations and trust. Build, buy or rebuild.
16:15–17:30	Exercise 2: value at risk LAB	Which revenue and cost lines depend on work that is getting cheap, and which on things that stay scarce.
DAY OUTPUT A task inventory, an exposure map and a short list of where the organisation's value is and isn't at risk.		

Day 2: From opportunities to a roadmap

Turn the exposure map into a funded, sequenced and governed plan.

Time	Module	Content and activity
09:00–10:30	Business models and pricing	Seat-based, usage-based and outcome-based pricing; what happens to margins when competitors automate; short-lived opportunities and how to test them cheaply.
10:30–10:45	Break	Coffee and questions.
10:45–12:15	Exercise 3: opportunity portfolio LAB	Score candidate initiatives by value, feasibility, risk and time to first evidence, and choose a sequence rather than a wish list.
12:15–13:15	Lunch	Lunch break.
13:15–14:00	Operating model and skills	Ownership, central versus embedded teams, upskilling the whole organisation, and measuring productivity without counting prompts.
14:00–15:00	Governance and risk	EU AI Act obligations and dates as they stand on the course date, ISO/IEC 42001, the NIST AI RMF, data leakage, vendor lock-in and over-automation.
15:00–15:15	Break	Coffee and questions.
15:15–16:15	Roadmap and business cases CAPSTONE	Each participant or team completes a 90-day roadmap with a one-page business case for every initiative.
16:15–17:30	Presentations and critique CAPSTONE	Roadmaps are presented and challenged by the group; each participant leaves with the three questions they still need to answer.
DAY OUTPUT A 90-day roadmap with owners, budgets, measurable outcomes and stop criteria, reviewed by the group.		

Labs and capstone

Lab 1: Exposure map	A task inventory of the participant's organisation, scored for exposure.
Lab 2: Value at risk	Revenue and cost lines linked to the tasks that are getting cheap.
Lab 3: Opportunity portfolio	Scored initiatives with a sequence and stop criteria.
Lab 4: Roadmap	A 90-day roadmap with owners, budgets, measures and governance notes.

Capstone specification

- A task inventory covering at least the organisation's three largest functions, scored for exposure.
- Three to five initiatives, each with an owner, a budget range, a measurable outcome and a stop criterion.
- A build, buy or rebuild decision for every initiative that needs software.
- A note on how each initiative affects pricing, margins or the customer relationship.
- An upskilling plan that names who is trained, in what, and how progress is checked.
- A governance page that maps the initiatives to the EU AI Act risk categories and the organisation's data-protection duties.

Tools and materials

Evidence	Benchmark reports and system cards (GDPval, METR, provider system cards) Primary sources, so participants can check claims instead of relying on summaries of them.
Exercises	Task-inventory and scoring templates The same scorecard Course 4 uses, so strategy and delivery speak one language.
Demonstrations	A coding agent and a desktop agent (for example Claude Code and Claude Cowork) Shows current capability on realistic work instead of describing it.
Economics	A cost model per task: people, tools, inference and review time Makes business cases comparable across initiatives.
Governance	EU AI Act, ISO/IEC 42001 and NIST AI RMF The frameworks a board, an auditor or a customer in Europe is most likely to ask about.

Discussed, not taught hands-on

- Model training, fine-tuning and infrastructure are explained only as far as a decision-maker needs them to judge cost and lock-in.
- Legal interpretation of the AI Act for a specific system needs counsel; the course prepares the questions to ask.
- Hands-on automation building is the subject of Course 4; teams that leave with a roadmap usually continue there.

The course helps leaders make and defend decisions about AI. It can't produce forecasts anyone could guarantee: capability is changing quickly, so the roadmap is built to be revised every quarter, and its business cases are estimates with stated assumptions.