From uncertainty to defensible decisions

Turning uncertain data — and AI-generated advice — into capital decisions your organisation can defend.

OptiField · MaestroOS™ | Five-day intensive workshop plus a follow-up session four weeks on · up to 20 participants
Delivered at your site or a host venue, independently or in partnership with a host organisation · taught on public field data

Course Overview

Every capital decision on a producing asset rests on numbers that are less certain than they appear. This course takes participants from where that uncertainty comes from — production and pressure data first, then wells, seismic and geological interpretation — through to the decisions it should drive: what to fund, what to defer, in what order, and what information is worth buying first.

What makes this course different

- **It ends in a real capital decision.** The week does not close on a summary of methods. Teams are given a budget, a costed menu of options and a deliberately ambiguous instruction from management, and must produce and defend an allocation — which they then have to revise when circumstances change mid-exercise.
- **Leadership takes part.** Asset managers and technical leadership join for the final day and question the teams directly. Technical staff practise defending a recommendation to the people who actually approve capital, and leadership sees how their own instructions are interpreted.
- **Participants apply it to their own asset, and come back.** Each day closes on a live decision from the participant's own portfolio, and the cohort reconvenes four weeks later to present what they did with it. The course is designed to change practice, not to be remembered fondly.

Learning Outcomes

On completion, participants will be able to:

1. Frame an asset decision and identify genuinely distinct alternatives, including defer and do nothing.
2. Distinguish reducible from irreducible uncertainty, and parameter uncertainty from structural uncertainty.
3. Propagate uncertainty through to production, recovery and economic outcomes, and interpret the resulting distribution rather than its mean.
4. Construct and evaluate a decision tree, including expected value, certain equivalent, sensitivity and the value of flexibility.
5. Calculate and interpret the value of information, and identify data that cannot change a decision.
6. Price the cost of delay, and use it together with duration to sequence approved work.
7. Evaluate a recommendation against risk appetite, capital constraints and corporate policy.
8. Critique an AI-generated recommendation against its evidence and provenance.
9. Present and defend a capital allocation recommendation under uncertainty, and state in advance what would change it.

Knowledge and Understanding

Participants will gain practical, in-depth knowledge in areas including:

- **Production and pressure data:** rates and allocation, metering error, well tests, production logs, pressure surveys — the primary evidence, and the most under-used

- **Wells and seismic:** logs, cores and petrophysics; where seismic and 4D genuinely help, and where the production history is already the stronger constraint
- **Representation:** distributions, ensembles, priors and scenarios — and what each can and cannot express
- **Propagation:** how subsurface uncertainty flows through to production, recovery and cash flow
- **Model limits:** structural uncertainty, data provenance, and the confidently wrong forecast
- **Decision analysis:** framing, alternatives, decision trees, expected value, the value of flexibility
- **Value of information:** when data acquisition is worth funding, and when it is not
- **Cost of delay:** what postponing a decision costs per unit time, and how to sequence what has been approved
- **Risk and portfolio:** risk appetite, correlated exposure, and identifying the robust core of a plan
- **Corporate policy:** how hurdle rates, screening criteria and governance change decisions independently of the data
- **AI decision support:** interrogating an AI assistant effectively, and recognising when to challenge it

Why Send Your Staff on This Course?

The course improves the quality of the recommendations that reach your investment decisions, building the analytical habits and shared language that make technical advice usable at the capital allocation table:

- **Improve decision quality, not only model quality** — the part your organisation controls
- **Get more from the data you already have** — before buying more
- **Reduce late surprises** — uncertainty surfaced while it can still change the plan
- **Direct data spend where it changes outcomes** — and stop funding studies that cannot
- **Decide when, not only what** — timing and sequencing priced rather than assumed
- **Bridge technical and investment decision-making** — one vocabulary, subsurface to gate
- **Build capability against recognised practice** — not a proprietary framework

Who Should Attend?

Reservoir, production and petroleum engineers, geoscientists, and the planning, commercial and economics staff who work alongside them. It is equally relevant for early-career technical staff building analytical foundations and for experienced personnel who recommend or approve capital. Asset managers and technical leadership join the cohort for the final day, when teams present and defend their capital allocation.

The material assumes working familiarity with production and pressure data from a producing asset. No specialist background in probability, statistics or decision analysis is required — those are built during the week.

Course Structure

The course is delivered in four parts:

1. **Pre-course briefing.** A short online session with the facilitators to set expectations and outline timelines. Course content is tailored to candidate needs, and where required to your own portfolio context, at no additional cost.
2. **Five-day workshop.** An intensive workshop for up to 20 participants, working in teams on a data-rich public field. Four technical days build the method; on the final day leadership joins for a capital allocation capstone under a live change in circumstances, which teams present and defend. Includes interactive sessions, group activities and participant-led teach-backs. Where group size exceeds 20, additional workshops can be arranged back-to-back. A three-day programme for those who approve rather than build technical work is available separately.
3. **Follow-up session, four weeks on.** A two-hour remote session in which participants bring the decision they have taken back to their own asset, present how they framed it, and receive peer and faculty feedback. This is where the method converts into practice rather than notes.
4. **Decision Impact Project (Optional).** A live decision on the participant's own asset, framed and analysed using the course method, approved by the line manager and presented to local management. Typically completed within six months, with an optional review at three months comparing what happened with what was expected.

Course Director

Milana Ayzenberg is a PhD geophysicist with an MSc in applied mathematics and 25 years in subsurface characterisation, uncertainty quantification and decision analysis. She taught Reservoir Geophysics at NTNU as Acting Professor.

At Equinor she conditioned reservoir models on dynamic data, then held portfolio responsibility across a $5 billion programme of field development plans and production optimisation, where the standing question was which work to do, and in what order.

She delivered the invited EAGE E-Lecture on seismic 4D inversion for automated history matching, has been Associate Editor of GEOPHYSICS and Geophysical Prospecting, and is co-founder and CTO of OptiField.

Delivery is supported by a second facilitator with operator-side experience, who runs the syndicate work and plays the management role in the Friday capstone.

The Learning Environment

INDEPENDENT METHOD, SUPPORTED BY SOFTWARE

The course teaches a method that stands on its own. Participants build a decision tree and value information by hand before they use any software, so that they understand what the tool is doing rather than trusting its output.

MaestroOS™ Light is then used as the interactive learning environment: participants explore uncertainty, build decision trees, run value-of-information calculations, and test what would have to change before a recommendation reverses — in minutes rather than hours, which is what makes the capstone possible in a single day.

What the course is not: it does not teach how to build AI models, run interpretation workflows or orchestrate agents. Those are well covered elsewhere and increasingly common. This course is about the decision those outputs feed — and about testing an AI-generated recommendation before capital is committed against it.

One session is spent learning to challenge the built-in AI assistant rather than defer to it. Teams are given a case where the assistant's confident answer is not well supported, and are scored on catching it and stating what evidence is missing. Participants return able to use AI decision support, and able to recognise when not to.

Progress Tracking

Full course administration is included and can be tailored to client requirements. Sponsoring management receives a participation and engagement summary after the workshop week, and again after the follow-up session. Where Decision Impact Projects are undertaken, a summary of the decisions addressed across the cohort can be provided — giving your organisation visibility of what the course actually changed, not only who attended.

Queries

Connect with enquiries@apes-energyevolution.com.