

Anatomy of a Fixed-Price AI R&D Phase: 250 Mandays, Five Roles, and the Discount Math

A fixed-price line item for an applied-AI research phase looks, to a buyer, like a single defensible number. It is not single and it is not a guess. This whitepaper reconstructs one such number from the engagement workbook: a 250-manday phase for a major operator in Oman, staffed by five roles on a fixed manday ladder (20 principal-QA, 60 senior researcher, 30 PM-plus-senior, 70 PhD researcher, 70 junior researcher), priced at academic day rates of 1600, 760, 760, 640 and 480 US dollars. Those academic rates run near half the open-market commercial ones (a principal at 1600 against 3200, a senior researcher at 760 against 1727), which is the single largest lever in the whole build and is available only because delivery routed through a university partnership. The role stack sums to 178,800 US dollars; a 35 percent bulk discount removes 62,580; a 7 percent contingency adds 8,135.40; and a 45,000 US dollar Phase-3 infrastructure package is carved out and priced separately, landing the phase at 169,355.40 US dollars. Behind the fixed price sits a two-party funding structure denominated in Omani rial: 165,000 OMR split 95,000 to the operator-facing firm and 70,000 to the university, each further split across a Phases 1-2 runway and the Phase-3 build. We argue that reading a fixed-price phase as an assembled artefact, rate arbitrage plus discount plus contingency plus carved-out compute, is the only way a buyer can tell a disciplined number from a hopeful one.

Tannistha Maiti

Senior AI Researcher

Narendra Patwardhan

Research Collaborator

FIXED-PRICE / AI-RD-PRICING / MANDAY-LADDER / DAY-RATE-ARBITRAGE /
ACADEMIC-PARTNERSHIP / BULK-DISCOUNT

CONTENTS

01	The ladder: five roles, 250 mandays, one fixed shape
02	The lever that halves the labour cost
03	From role stack to fixed price: the waterfall
04	Why the sequence is the substance
05	Behind the fixed price: a two-party budget in rial
06	What the fixed price does and does not commit
07	How a buyer should read a fixed-price phase
08	What the anatomy is for
09	Limitations
10	References

A fixed-price line item for a research phase is a strange object. To the buyer it arrives as one number on one line, and the whole apparatus that produced it, the staffing, the rate card, the discounting, the contingency reserve, the compute that was quietly moved off the labour line, is compressed out of view. What is left is a figure that either looks reasonable or does not, with no obvious way to interrogate the difference. For a phase we delivered to a major operator in Oman, that figure was 169,355.40 US dollars, and this paper is about the machinery behind the decimal.

We are writing as the people who built the technical work and sat close to how it was costed. The published record of this engagement already covers the science in depth, the set-prediction fracture detector, the vug quantification pipeline, the well-to-well correlation, and it covers the operating side too, the monthly burn rate of the research team and the marginal compute cost of digitising one more log. What it has never covered is how a single phase was commercially engineered: the manday ladder, the day-rate arbitrage, the discount and the carve-out that turn a stack of roles into a number a procurement team can sign. That is the gap this whitepaper fills. If you are the person on the buy side deciding whether a quoted phase is disciplined or hopeful, this is the anatomy you need.

The thesis is simple and we will hold it through every section: a 250-manday fixed-price phase is engineered, not quoted. An academic-versus-commercial rate ladder near half price, a 35 percent bulk discount, and carved-out infrastructure are what turn a role stack into a defensible number. Read the number as an assembled artefact and its logic becomes legible. Read it as a guess and you cannot tell a good one from a bad one.

The ladder: five roles, 250 mandays, one fixed shape

The phase was staffed by five roles, and the effort behind each was fixed in mandays before any rate was applied. This ordering matters. The team decided how much of each seniority the work required, then priced that effort, rather than starting from a budget and reverse-fitting the people. The manday ladder was 20, 60, 30, 70, and 70, summing to exactly 250 mandays.

Reading top to bottom: a principal, carrying quality assurance and technical oversight, at 20 mandays. A senior AI researcher and domain expert at 60. A project manager who doubled as a senior researcher at 30. A PhD-candidate researcher at 70. A junior researcher, also at 70. The shape of that ladder is itself a claim about the work. The bulk of the effort, 140 of

the 250 mandays, sits with the two most junior roles, which is what you expect from a build phase where a great deal of the labour is data preparation, experiment running, and iteration. The principal's 20 mandays are deliberately scarce: senior oversight is expensive and is spent on review and direction, not on production.

There is a second signal in the ladder that a buyer should learn to read: the presence of a dedicated quality-assurance role at the top. The principal's 20 mandays are not general management overhead. They are a specific, budgeted allocation for review of the research output, and pricing them as a distinct line rather than burying them in a blended rate is a statement that the work will be checked by someone senior and that the buyer is paying for that check explicitly. In a research phase, where the deliverable is a model and a set of results rather than a fixed artefact, that oversight line is often the difference between a result you can defend and one you cannot. A ladder with no senior QA time is cheaper on paper and more expensive in practice.


Look also at what the ladder does *not* contain. The two infrastructure and software-development roles that appear elsewhere in the workbook, the MLOps infrastructure developer and the software developer, are budgeted at zero mandays in this particular sheet. That is not an oversight. It is the labour-side expression of the same carve-out we will meet again below: the infrastructure work is priced separately, as a package, rather than as days on the labour ladder. Seeing a role at zero mandays on the labour sheet and a matching cost on the infrastructure line is exactly the kind of internal consistency a buyer wants to find, because it shows the two lines were reconciled rather than double-counted.

Every one of those manday counts is a sourced figure from the phase effort-calculation workbook, not an estimate we are reconstructing after the fact. The same is true of the day rates, which is where the number starts to become interesting.

The lever that halves the labour cost

The five roles were priced at academic day rates: 1600, 760, 760, 640, and 480 US dollars per day, matched to the ladder above. Multiply rate by mandays and the role stack sums to 178,800 US dollars. That is the gross labour cost of the phase before any commercial adjustment, and it is the base of the waterfall.

The important fact about those rates is what they are *not*. They are academic rates, available because the engagement was delivered through a university partnership, and they run near half the open-market commercial rates for the same seniority. The reference rate table in the workbook makes the gap explicit. A principal or CEO-level day costs 3200 commercially and 1600 academically. A senior researcher costs 1727 commercially and 760 academically. A PhD-level researcher costs 1200 against 640, a junior 640 against 480. Priced at those commercial rates, the identical 250-manday ladder would have summed to 348,230 US dollars rather than 178,800. The university route removes roughly 170,000 US dollars from the labour line before a single discount is applied.



The day-rate arbitrage that anchors the fixed price. Each role is drawn as a paired bar, the faint outline being the open-market commercial rate and the solid teal being the academic rate offered through the university partnership. The academic rate lands near half the commercial one at every rung: principal 1600 against 3200, senior researcher 760 against 1727, PhD researcher 640 against 1200, junior researcher 480 against 640. The reading lever switches between raw day rate and booked cost, which multiplies each rate by the role's contracted mandays, so the arbitrage reads both as a per-day gap and as dollars removed from the stack. The orange bracket is the only element that argues: it marks the rung with the widest absolute gap. Every rate and manday count is sourced from the engagement's effort-calculation workbook; none is illustrative.

This is the single largest lever in the entire build, and it is worth being precise about why it exists. It is not a negotiation win and it is not a volume effect. It is a structural feature of who delivers the work. A university partner can offer research labour at academic rates because the institution's economics, its overheads, its funding, its purpose, are different from a commercial consultancy's. Routing the phase through that partnership is a decision made at the shape of the engagement, long before the price is quoted, and it is why the two rate columns in the arbitrage ladder above sit so far apart at every rung. A buyer who sees only the academic column sees a number that looks cheap; a buyer who sees both columns understands that the cheapness is engineered into the delivery structure, and can ask the right question, which is whether that structure will hold for the life of the work.

The Omanization dimension sits directly on top of this. A university partnership in Oman is not only a rate mechanism; it is a training and national-capability mechanism. Across the wider programme, dozens of young Arab professionals and Omani nationals were trained through the partner institution, which means the academic rate is buying two things at once: research labour at half the commercial cost, and a pipeline of local talent that a purely commercial delivery would not produce. For a national operator, that second product is not a side effect. It is often the reason the engagement is structured through a university in the first place.

From role stack to fixed price: the waterfall

With the labour base set at 178,800 US dollars, the phase price is built in three more moves, and the order of those moves is the argument of this whole paper.

First, a 35 percent bulk discount removes 62,580 US dollars, taking the stack down to 116,220. This is the commercial gesture that acknowledges the size and the multi-phase commitment of the work: a buyer taking a whole phase, not a day here and there, is priced accordingly. Second, a 7 percent contingency adds 8,135.40 US dollars, a reserve against the ordinary overruns of research work, taking the running figure to 124,355.40. Third, and separately, a 45,000 US dollar Phase-3 infrastructure package is added on top, landing the phase at 169,355.40 US dollars.



A 250-manday fixed-price AI R&D phase read as an engineered number rather than a quoted one. The left column prices five roles from a fixed manday ladder (20/60/30/70/70) at academic day rates (1600/760/760/640/480 USD), summing to an 178,800 USD role stack. The waterfall on the right then builds the price step by step: down by the 35 percent bulk discount (the orange step, the one element that argues the thesis), up by the 7 percent contingency, and up again by the 45,000 USD Phase-3 infrastructure slab that was carved out and priced separately, landing on the 169,355.40 USD fixed price. The day-rate lever reprices the same stack at commercial rates, which run near double the academic ones (principal 3200 vs 1600, senior researcher 1727 vs 760 per day), so the reader sees the rate arbitrage that anchors the whole number. Every value plotted is sourced from the engagement's effort-calculation workbook and infrastructure-cost memo; none is illustrative.

The waterfall above is the spine of the argument, so read it deliberately. The tall teal bar on the left is the academic role stack. The orange step is the 35 percent bulk discount, the one move that visibly pulls the number down, and it is orange because it is the element that most changes the buyer's perception of the price. The two small teal steps that follow are the contingency uplift and the carved-out infrastructure. The final bar is the fixed price. Toggle the day-rate lever and the whole structure reprices from the commercial column, and the point of the thesis becomes visible in one gesture: without the academic arbitrage, the same ladder, same discount, same contingency, same infrastructure, produces a materially larger number. The discount is the move buyers notice; the rate arbitrage is the move that actually sets the scale.

Two design choices in that build deserve to be named, because they are the difference between a disciplined number and a round one.

The first is that infrastructure is carved out. The 45,000 US dollars of Phase-3 compute and tooling is not folded into the labour rates and it is not inside the discount base. It sits on its own line, added after the discount and contingency are computed on labour alone. This matters for honesty and for control. Compute is a real, separable cost with its own drivers, and a buyer who can see it on its own line can question it on its own terms, can ask what hardware, what duration, what utilisation. Blend it into a day rate and it becomes invisible and unaccountable. The published operating-side analysis of this engagement makes the same point from the run-cost angle: the marginal cost of the next digitised log is a serving property of the infrastructure, not a training cost, which is exactly why compute belongs on its own line rather than smeared across labour.

The second is that the discount is applied to the labour base only, before contingency and before infrastructure. That sequencing is not neutral. A 35 percent discount computed on a base that included the 45,000 US dollars of infrastructure would have removed nearly 16,000 US dollars more, and would have discounted a hardware cost that has no business being discounted as if it were negotiable labour. Applying the discount to labour, then adding contingency on the discounted labour, then adding infrastructure at full cost, is the sequence that keeps each component priced on its own logic.

A skeptical buyer might ask whether the 35 percent is real or theatrical, a headline concession off an inflated list price. The arbitrage ladder answers that question. The discount is taken off the academic base of 178,800 US dollars, which is already the low column, not off the commercial base of 348,230 that the same work would have carried on the open market. A vendor engineering a fake discount would have done the opposite: quote at or near commercial rates, then apply a large-looking percentage to arrive back at a merely ordinary number. Here the cheap rate and the discount stack in the buyer's favour, one on top of the other, which is the signature of a genuine concession rather than a pricing optic. The 35 percent is a real 62,580 US dollars removed from an already-discounted labour line.

Why the sequence is the substance

It is tempting to treat discount, contingency, and carve-out as arithmetic details that wash out in the total. They do not. Change the order and you change both the number and what the number means.

Consider the contingency. Seven percent on the discounted labour base of 116,220 US dollars is 8,135.40. Seven percent computed on the *undiscounted* base of 178,800 would have been 12,516, a reserve nearly 4,400 US dollars larger, and a reserve sized against a labour cost the buyer is not actually paying. Contingency is meant to cover the risk of the work that will be done at the price that will be charged, so it belongs on the discounted base. The workbook computes it there. That is a small decision with a clear rationale, and it is the kind of decision that distinguishes a costed number from a marked-up one.

Or consider what the carve-out protects. If infrastructure were inside the discount base, a larger discount would look like generosity while quietly eroding the compute budget the work actually needs, and the first casualty of an under-funded compute line is the science. By pricing the 45,000 US dollars separately and at full value, the structure guarantees that the discount is a commercial concession on labour and not a stealth cut to the hardware that the fracture and vug models have to run on. The buyer gets a real discount; the delivery keeps its real compute.

THE PHASE PRICE, ASSEMBLED COMPONENT BY COMPONENT

[Formula](#)[LaTeX](#)[Copy LaTeX](#)

$$P = \underbrace{\sum_i d_i r_i}_{\text{role stack}} (1 - \delta) (1 + \gamma) + I$$

The expression above states the whole build in one line. The role stack is the sum over roles of mandays times day rate. The discount factor removes 35 percent, so it multiplies by one minus delta. The contingency factor adds 7 percent to what remains, so it multiplies by one plus gamma. And the infrastructure term is added outside both factors, because it is neither discounted nor contingency-loaded. Substitute the sourced values, a stack of

178,800, delta of 0.35, gamma of 0.07, and infrastructure of 45,000, and the expression returns 169,355.40. There is nothing hidden in the total; every term is nameable and every term traces to the workbook.

Behind the fixed price: a two-party budget in rial

The 169,355.40 US dollar figure is the price of one phase to the buyer. Sitting behind it is a funding structure denominated in Omani rial that describes how the two delivery parties, the operator-facing AI firm and the university, split the wider programme. The programme budget was 165,000 OMR, and the split was 95,000 to the firm and 70,000 to the university.

Each party's share breaks further across the phases. The firm's 95,000 OMR is 39,000 across the Phases 1-2 runway and 56,000 in Phase 3. The university's 70,000 OMR is 45,000 across Phases 1-2 and 25,000 in Phase 3. Read by party, the firm carries the larger share, 58 percent of the budget. Read by phase, the weight shifts toward Phase 3, the build, where the firm's share in particular concentrates.



The 165,000 OMR programme budget read as a split between the two delivery parties rather than as a single fixed price. Each stacked bar breaks a party's total into its Phases 1-2 runway and its Phase-3 build: the operator-facing firm carried 95,000 OMR (39,000 across Phases 1-2 and 56,000 in Phase 3), and the university partner carried 70,000 OMR (45,000 then 25,000). The split lever re-buckets the same total by phase instead of by party, so the reader sees both who carried the money and when it was committed. The orange slab is the only element that argues: the firm's Phase-3 bucket, the single largest, which is where funding concentrated once the 250-manday build phase this paper anatomises began. Every figure is sourced from the engagement's effort-calculation workbook funding sheet; none is illustrative.

The bridge above shows why this matters for a buyer even though the buyer sees only the fixed price. The funding split tells you where the work, and the risk, actually sits. The university's larger relative share falls in the early runway phases, which is consistent with its role in dataset development, training, and the academic-rate labour that anchors the whole cost structure. The firm's share concentrates in the Phase-3 build, the orange slab, which is consistent with the engineering-heavy, infrastructure-heavy character of the phase this paper anatomises. A two-party structure is not an accounting curiosity; it is the reason the academic rate exists at all, because one of the two parties *is* the academic institution. The rial split and the dollar arbitrage are the same fact seen from two sides.

There is a governance point folded into this too. When two parties fund and deliver a programme jointly, the intellectual property, the publication rights, and the effort-versus-result obligations have to be negotiated at the structural level, not the invoice level. In this

engagement those terms were settled in the founding agreements, with research services explicitly framed as an effort obligation rather than a guarantee of a specific result, which is the correct framing for a research phase and one a buyer should expect to see. The commercial structure and the legal structure are built together, and the fixed price is the visible tip of both.

What the fixed price does and does not commit

A fixed price on a research phase is a specific kind of promise, and it is worth being exact about which risks it transfers and which it leaves with the buyer. This is where the effort-versus-result framing does real work rather than sitting in the contract as boilerplate.

The manday ladder fixes effort. The vendor is committing to deliver 250 mandays of the specified seniority mix at the agreed rates, and the fixed price converts that commitment into a single number the buyer can budget against. What the fixed price does not do, and honestly cannot do, is guarantee a particular research outcome. The nature of a research phase is that the result is discovered, not manufactured to spec, and a vendor who fixed the price *and* guaranteed a specific accuracy figure would be either misunderstanding the work or mispricing the risk. Framing the services as an effort obligation is the correct and honest position, and it is why the manday ladder, not a deliverable specification, is the object being priced.

This has a direct consequence for how contingency should be read. The 7 percent reserve is not a hedge against the research failing to reach a target, because no such target is being underwritten. It is a hedge against the ordinary friction of the work: an experiment that has to be rerun, a data intake that arrives dirty and needs a cleaning pass that was not scoped, a week where the compute misbehaves. Sized at 8,135.40 US dollars on the discounted labour base, it is a modest, proportionate reserve for those frictions, not a margin buffer dressed up as risk cover. A contingency several multiples larger than this, on a phase of this size, would be a signal that the vendor is either uncertain about the scope or padding the number.

The multi-phase structure changes the risk calculus again. Because the programme is staged, the buyer is not committing to the entire body of work at the point of signing one phase. Each phase is a decision point, and the funding split we just traced, with its Phases 1-2 runway distinct from the Phase-3 build, is the financial expression of that staging. The

infrastructure carve-out fits the same logic: a compute package sized for the build phase is committed when the build phase is, not amortised blindly across the whole programme. Staging is a risk-management feature for the buyer as much as a delivery convenience for the vendor, and a fixed-price phase inside a staged programme is a much safer object to sign than a fixed price for an entire multi-year research effort quoted up front.

How a buyer should read a fixed-price phase

The reason to reconstruct one number in this much detail is that the method generalises. A buyer handed a fixed-price research phase cannot see the workbook, but can ask the questions that would expose its shape, and the answers separate a disciplined vendor from a hopeful one.

- Is the price built from a manday ladder, or is it a round number? A ladder that sums to an exact figure like 250 mandays, with effort assigned per role before rates are applied, is evidence the work was scoped before it was priced. A suspiciously round total is evidence of the reverse.
- What rate basis is the labour priced at, and why? If the rates are low, the buyer should understand the mechanism, an academic partnership, a national-capability programme, an offshore team, rather than assuming the low number will hold. A rate that is cheap for a structural reason is durable; a rate that is cheap for no stated reason is a future change order.
- Is the discount applied to labour alone, and in what order? A discount computed on a base that includes pass-through costs like infrastructure is either a mistake or a tactic. The sequence should be legible.
- Is infrastructure on its own line? Compute is separable, has its own drivers, and should be questionable on its own terms. If it is invisible inside a blended rate, the buyer has lost the ability to control the single most volatile cost in an AI programme.
- Is contingency sized against the price that will actually be charged? A reserve computed on an undiscounted base is oversized and mispriced.

None of these questions require access to the vendor's spreadsheet. They require only the assumption that a fixed price is an assembled artefact, and the willingness to ask how it was assembled. The 169,355.40 US dollar phase in this paper answers all five, which is the whole reason it reads as a disciplined number rather than a hopeful one.

What the anatomy is for

We opened with the observation that a fixed-price line item compresses its own machinery out of view. The purpose of this paper has been to decompress one such number completely: five roles, 250 mandays, academic rates at roughly half commercial, a 35 percent discount on labour, a 7 percent contingency on the discounted base, and 45,000 US dollars of infrastructure carved out and added at full cost, resolving to 169,355.40 US dollars, and sitting inside a 165,000 OMR two-party budget split 95,000 to the firm and 70,000 to the university.

The number is specific to this engagement. The method is not. Every applied-AI research phase can be read the same way, as a rate basis times an effort ladder, adjusted by a discount and a contingency whose bases and order matter, plus a compute line that should stand on its own. A vendor who can show that structure is telling you the number was engineered. A vendor who cannot is asking you to trust a total. The thesis holds: a fixed-price phase is built, not guessed, and the buyer who knows how it was built is the buyer who can tell the difference.

Limitations

This paper reconstructs the commercial anatomy of a single research phase for a single operator, and its numbers should be read with that scope in mind. The manday ladder, day rates, discount, contingency, and infrastructure figure are the values recorded in the phase effort-calculation workbook and the infrastructure-package cost memo for this engagement; they are not a market benchmark, and other engagements will differ in every parameter. The academic-versus-commercial rate arbitrage is specific to delivery through a particular university partnership in Oman and to the seniority mix of this phase; the roughly one-half ratio should not be assumed to generalise to other institutions, regions, or role ladders. The Omani rial funding split describes the wider programme's two-party structure, not a line-item mapping to the fixed-price phase, and the two are denominated in different currencies and set at different times, so they should not be reconciled arithmetically. The instruments in this paper plot only sourced figures and carry no illustrative inputs, but they are a reading aid, not an audited financial statement. Finally, this is an analysis of how the phase was priced, not a claim about whether the price was correct; judging value would require the delivered outcomes, which the technical record covers separately.

References

Effort-calculation workbook (engagement proposal artefact), Phase 3. The primary source for the manday ladder (20/60/30/70/70), the academic day rates (1600/760/760/640/480), the 35 percent bulk discount, the 7 percent contingency, the carved-out infrastructure figure, the commercial-versus-academic reference rate table, and the two-party OMR funding split. Internal, held by the delivery team.

Infrastructure-package cost memo (engagement proposal artefact), Phase 3. The source for the 45,000 US dollar Phase-3 infrastructure package priced separately from labour. Internal, held by the delivery team.

Founding service and collaboration agreements (engagement governance artefacts). The source for the two-party delivery structure, the effort-versus-result framing of the research services, and the intellectual-property and publication terms referenced in the funding-split discussion. Internal, held by the delivery team.

Anatomy of a Fixed-Price AI R&D Phase: 250 Mandays, Five Roles, and the Discount Math

Published February 2026. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/anatomy-of-a-fixed-price-ai-rd-phase-250-mandays-five-roles-discount-math>

Copyright EarthScan. All rights reserved.