ANNEX A · KENYA INSTITUTE FOR CLINICAL AI

A

The Business Case

Kenya Institute for Clinical Artificial Intelligence — a proposal to the Board

Dr Neal Aggarwal

Written in a personal capacity

AUGUST 2026

Annex A to Another Arrow in the Quiver: What I Would Build to Teach Clinical Artificial Intelligence to Kenya's Doctors, Surgeons and Nurses.

I write here in a personal capacity.

1. The decision I am asking the Board to take

I am asking for three things, in sequence, with a gate between each.

Decision 1 — Approve Phase 0. Authorise USD 1.94 m over months 0–12 to establish the Institute: nine founding posts, the simulation centre and sandbox build, and negotiation of competency standards with the two councils. This is the only decision that commits capital.

Decision 2 — At month 15, review the pilot. An independent evaluation of a 120-learner, single-hospital pilot will be published. If it shows the training does not change workplace behaviour, the correct action is to redesign, not to scale. I am asking the Board to hold me to that.

Decision 3 — At month 15, approve or decline Phase 2. Scaling commits the recurrent cost base to roughly USD 3.6 m a year rising to USD 4.9 m.

I would rather the Board declined at gate 2 than approved a programme that scaled an unevaluated pilot. The commonest way institutions of this kind fail is precisely there.

2. Executive summary

Five-year expenditure	USD 17.37 m (KES 2,260 m)
Five-year earned income	USD 5.34 m (KES 694 m)
Five-year subsidy required	USD 12.03 m (KES 1,566 m)
Learners certificated over five years	21,746
Blended cost per learner certificated	USD 799
Subsidy per learner	USD 554
Earned-income cover, Year 5	49%
Core establishment at steady state	71 FTE

The honest headline: this does not wash its own face, and it is not designed to. At the fee levels that Kenyan clinicians can actually pay, earned income reaches 49% of cost by Year 5. Full cost recovery would require Year 5 income of USD 4.92 m — 2.1 times what the model projects. The Institute is a subsidised public good with a substantial and growing earned-income component, not a business.

That is the argument the Board has to accept or reject. Everything else in this document is detail.

What the Board gets for the subsidy is set out in §9: a research asset, a regional teaching franchise, a defensible claim to have built the first institution of its kind anywhere, and — the reason I care most — a measurable reduction in the rate at which AI contributes to harm in this hospital's own wards.

3. Why this institution, and why here

I will not restate the clinical argument; it is in the blueprint. Three points bear on the Board's decision specifically.

The host gets the competence first. Whatever the national ambition, the first cohorts are this hospital's own staff. The pilot is single-site by design. If the training works, this institution's clinicians are the first to hold the certificate and the first to have their AI-assisted practice audited.

The alternative is not "no cost". Clinicians here are already using these systems, untaught and unassessed. The cost of that is currently invisible because nobody is measuring it — which is not the same as zero. One incident in which AI contributes to a serious adverse event, with no training record and no documented governance, would cost more in inquiry and reputational terms than a year of this programme.

No teaching hospital in the region has moved. The window in which "first" is available is short. The AAMC reports North American medical schools going from 53% to 77% AI-curricular adoption in a single year. That curve arrives here.

4. What is being built, in one page

A semi-autonomous training and standards body hosted by this institution, with its own governing board and budget line, doing five things: teaching, assessment and certification, simulation and sandbox, evaluation and research, and standards and advisory work.

The teaching runs across **five professional tracks** — physicians; surgeons and proceduralists; nursing and midwifery; administrators and clinical managers; health informaticians — and **five gated levels** from a twelve-hour common core to a twelve-month fellowship. Progression requires a passed assessment and a countersigned workplace logbook. There is no time-served route.

The pedagogical core is the trained detection of error in AI output. A 2025 randomised trial found that physicians who had already completed twenty hours of AI-literacy training still deferred to deliberately erroneous model output. That finding is why the cur-

riculum is built around discernment rather than familiarity, and why error-catch rate is a conjunctive pass requirement.

5. How the model was built

Every figure in this document is computed, not asserted. The model lives in `financial_model.py` alongside this text. Change an assumption, re-run, and the tables regenerate. The Board should treat the assumptions, not the outputs, as the thing to argue with.

5.1 Assumption register

Assumption	Value	Basis and confidence
Currency	USD primary	Board preference; KES shown at planning rate
Exchange rate	KES 130 = USD 1	Mid-market, August 2026. **Revisit quarterly**
Employer on-costs	26% on gross	Statutory contributions, pension, medical cover. **Verify with HR**
Pay and non-pay inflation	5% p.a. from Y2	Planning figure. **Low confidence**
Salary bands	See §5.2	**Lowest-confidence input in the model.** See the caution below
FTE ramp	16 / 34 / 60 / 71 / 71	From the blueprint's four phases
Fee levels	USD 45 / 220 / 480 / 650 / 6,000	Judgement. Tested at half rate in §7
Throughput	See §6.2	Judgement, gated on pilot
Contract income	0 / 180k / 520k / 900k / 1.25 m	**Least reliable income line.** Tested to zero in §7

A caution on the salary bands. I attempted to benchmark these against published Kenyan salary data and stopped, because the public aggregators are self-reported and internally incoherent — one widely-cited source lists specialist surgeons at KES 181–217k per month and general surgeons at KES 694–800k in the same table. Building a five-year cost base on that would be false precision.

The bands below are therefore **stated planning assumptions**, set at what I judge would be required to recruit and retain each post in Nairobi in competition with private hospitals and the technology sector. They are the single input most likely to be wrong. §7 tests the model at ±20% on the whole payroll. The Board's HR function should replace them with real scales before Decision 1 is taken.

5.2 Establishment and payroll

Unit	Posts	Gross USD	Notes
Executive	1	96,000	Director / Chief Executive
Office of Quality and Evaluation	6	200,000	Reports to the Board, not the CEO
Ethics and Data Governance	3	110,000	Incl. certified data protection officer
Curriculum and Pedagogy	15	650,000	Director, 4 designers, 6 clinical leads, psychometrician, 3 editors
Faculty and Delivery	18	700,000	5 track leads, 10 certified instructors, 2 sim faculty, manager
Engineering and Platform	13	522,000	Competing directly with the private tech market
Simulation and Clinical Labs	7	196,000	
Operations and Registry	8	194,000	Registry must be auditable and defensible
Total	**71**	**2,668,000**	
Fully loaded at 26% on-costs		**3,361,680**	Average USD 47,348 per post

Two posts are worth the Board's specific attention.

The assessment psychometrician (USD 60,000). This is a scarce skill in the region and the temptation will be to defer the post and let a clinician "handle assessment". That decision would hollow out the certification before anyone noticed, and the Institute's entire value rests on the certificate meaning something.

The engineering unit (13 posts, USD 522,000). These posts compete with the private technology sector, not with the health sector. Pricing them on hospital scales will not fill them.

5.3 Capital

Item	USD	Year
Simulation centre fit-out — 2 consultation rooms, ward bay, theatre, debrief	480,000	1
Sandbox, platform build and de-identified case corpus	350,000	1
Headquarters office fit-out	110,000	1
Laptop, AV and recording fleet	90,000	1
Assessment hall, 30 invigilated stations	120,000	2
Mobile delivery unit — vehicle, generator, satellite uplink, kit	140,000	2
Total capital	**1,290,000**	

The mobile unit is the line most likely to be cut, and cutting it would be a mistake. A substantial fraction of the workforce the Institute most needs to reach works at Level 3 and Level 4 facilities they cannot leave for a week. Without it, the Institute becomes a thing that happens in Nairobi to people who can afford to come to Nairobi — failing at its actual purpose while appearing to succeed.

5.4 Recurrent non-pay at steady state

Item	USD/yr
Fellowship stipends — 8 fellows, from Year 4	192,000
Cloud, compute and model API access	120,000
Utilities, maintenance and premises	90,000
Travel and mobile-unit running costs	80,000
Clinical champions — sessional, 0.2 FTE each	72,000
Standardised patients — sessional	60,000
Accreditation, legal, audit and insurance	45,000
External examiner and triennial peer review	30,000
Total at steady state	**689,000**

Ramped at 30% / 55% / 85% / 100% / 100% across the five years.

6. Income

6.1 Fees

Level	Award	Contact	Fee USD
L1 Foundation	Certificate in Clinical AI Foundations	12 h	45
L2 Practitioner	Practitioner Certificate (by track)	36–40 h + logbook	220
L3 Advanced	Advanced Certificate (specialty named)	60 h + project	480
L4 Faculty	Certified Instructor	pedagogy core + practicum	650
L5 Fellow	Clinical AI Fellowship	12 months	6,000

Fees are set at what I judge a Kenyan clinician will actually pay from their own pocket for CPD that their council requires. **They are deliberately not set at cost recovery.** A Foundation certificate priced to recover its true cost would exclude precisely the cadres — nursing, clinical officers at county facilities — where the yield is highest.

The Fellowship fee is nominal; in practice fellows will be sponsored or bursaried, and the model does not depend on fellowship income.

6.2 Throughput

Year	L1	L2	L3	L4	L5	Total
1	120	—	—	—	—	120
2	1,200	150	—	40	—	1,390
3	3,500	600	120	60	—	4,280
4	5,500	1,100	240	60	8	6,908
5	7,000	1,600	360	80	8	9,048

Year 1 is the pilot and is free at the point of use — the Institute cannot in conscience charge for training it has not yet shown to work.

6.3 Institutional contracts and advisory

Year	1	2	3	4	5
USD	0	180,000	520,000	900,000	1,250,000

Hospitals buying cohort places for their own staff, plus paid evaluation and deployment-governance work for facilities procuring clinical AI systems. **This is the least reliable line in the model** and §7 tests it at zero.

There is a governance point here the Board should see clearly. Advisory income creates a conflict: an Institute paid to advise on a deployment cannot be a disinterested evaluator of it. The mitigation is in the founding instruments — the right to publish evaluation findings regardless of outcome is non-negotiable, and any contract that will not accept it is refused. **The Board should expect this to cost us contracts, and should want it to.**

7. The five-year position

USD	Y1	Y2	Y3	Y4	Y5
Expenditure					
Payroll incl. on-costs	757,562	1,690,310	3,132,044	3,891,565	4,086,143
Recurrent non-pay	149,100	287,018	465,751	797,604	837,484
Capital	1,030,000	260,000	—	—	—
Total expenditure	**1,936,662**	**2,237,327**	**3,597,795**	**4,689,168**	**4,923,627**
Income					
Course fees	—	118,650	425,675	800,729	1,142,333
Contracts and advisory	—	180,000	520,000	900,000	1,250,000
Total earned income	**—**	**298,650**	**945,675**	**1,700,729**	**2,392,333**
Subsidy required	**1,936,662**	**1,938,677**	**2,652,120**	**2,988,439**	**2,531,294**
Earned-income cover	0%	13%	26%	36%	49%

7.1 Unit cost

	Y1	Y2	Y3	Y4	Y5
Learners certificated	120	1,390	4,280	6,908	9,048
All-in cost per learner (USD)	16,139	1,610	841	679	544

The Year 1 figure is not a meaningful unit cost — it is the whole capital build divided by a pilot cohort. The number that matters is the **Year 5 marginal position: USD 544 per clinician certificated, all-in**, falling as throughput rises against a largely fixed cost base.

For comparison, that is in the region of a single day of specialist locum cover.

7.2 Sensitivity — Year 5 subsidy requirement

Scenario	Y5 subsidy (USD)	Δ vs base
Base case	**2,531,294**	—
Payroll bands 20% lower than assumed	1,714,065	−817,229
Throughput 30% above plan	2,188,594	−342,700
Throughput 30% below plan	2,873,994	+342,700
Fees held at half the assumed rate	3,102,460	+571,166
Payroll bands 20% higher than assumed	3,348,523	+817,229
Contracts fail to materialise entirely	3,781,294	+1,250,000

Read the sensitivity table before the base case. The two largest swings are payroll bands — the input I have least confidence in — and contract income, the line most likely to disappoint. Together they span roughly USD 2.1 m of Year 5 subsidy. Any Board discussion that treats the base case as a forecast rather than a midpoint is having the wrong conversation.

7.3 Break-even

Full cost recovery in Year 5 requires earned income of **USD 4.92 m**, against a modelled **USD 2.39 m** — a factor of **2.1**.

Getting there would require roughly doubling fees, which would price out the cadres the Institute exists to reach, or roughly doubling throughput, which the pipeline does not support at that quality bar. **I am not proposing either.** I am proposing that the Board fund a subsidised institution with a rising earned-income share, and hold it to publishing that share annually.

8. Financing

The subsidy requirement is USD 12.03 m over five years. I would not seek it from a single source.

Source	Indicative share	Comment
Host institution (cash and in-kind)	25%	Premises, finance, HR, legal, clinical access
Philanthropic and development partners	45%	Global health funders; the evidence case is strong and Kenyan
Council- and employer-funded CPD	20%	Grows as council recognition beds in
Research grants	10%	The evaluation programme is fundable in its own right

In-kind host contribution is real money and should be recognised as such: premises alone, at the floor area the simulation centre needs, is a material line.

The financing risk the Board should press me on is the 45%. Development-partner appetite for health-workforce training in Kenya is real but crowded, and no funder has been approached. That conversation cannot begin credibly until Decision 1 has been taken and the founding team exists.

---

9. What the host institution gets

Clinical safety in its own wards, measured. The Institute audits AI-assisted practice at three and twelve months post-training, with chart audit and workplace observation. No teaching hospital in the region currently has any measurement of this at all.

A research asset of unusual value. The de-identified Kenyan case corpus, and the sandbox interaction logs, constitute a dataset on clinician–AI interaction in an African clinical context that does not exist anywhere. It is publishable, fundable, and it is the kind of asset that attracts collaborators rather than requiring them to be recruited.

A regional teaching franchise. Phase 3 opens East African Community faculty exchange. The realistic ambition is that this hospital becomes where the region's clinicians come to be certified.

Recruitment and retention. Formidable clinicians increasingly ask what an institution will teach them. A funded fellowship and a certified instructor track are a recruitment argument.

Priority in the accreditation queue. The Institute will hold KMPDC and Nursing Council CPD provider accreditation. The host's own staff meet their annual CPD obligations in-house.

A defensible claim to primacy. No permanent national institution of this kind exists anywhere. If it is built here it will be studied, visited and copied.

---

10. Risks to this business case

Distinct from the programme risks in the blueprint, these are the risks to the *money*.

Risk	Likelihood	Impact	Mitigation
Salary bands materially understated	**High**	High	HR to replace bands before Decision 1; model re-run; ±20% already tested
Contract income disappoints	Medium	High	Tested to zero; not required for Phases 0–1
Development-partner funding does not close	Medium	High	Phase gates mean Phase 2 is not committed until funding is visible
Fees prove unpayable in the county cadres	Medium	Medium	Bursary line; employer-funded places; sliding scale by facility level
Engineering posts cannot be filled at these rates	**High**	Medium	Benchmark against tech sector, not health; budget for contractors in Y1
Capital overrun on the simulation centre	Medium	Medium	Fixed-price fit-out; phase the theatre room to Y2 if needed
Host absorbs cost without the independence guarantees	Low	**Severe**	Independence written into founding instruments as a condition of Decision 1

That last row is the one I would ask the Board not to trade away. An Institute whose evaluations can be suppressed by its host is worth less than nothing — it would launder deployments rather than test them.

11. What I am asking for at Decision 1

USD 1,936,662 for months 0–12, comprising USD 906,662 recurrent and USD 1,030,000 capital.

That buys: the governing board constituted and the independence instruments executed; the founding nine recruited; competency standards drafted and taken to KMPDC and the Nursing Council; the twelve-hour common core written, peer-reviewed and piloted; the sandbox and de-identified case corpus built; the first instructor cohort certified; and a 120-learner pilot at this hospital with an independent evaluation published at month 15.

At which point the Board decides whether any of this works.

11.1 A declaration, and a proposal on my own remuneration

I am a candidate for the Director post. The Board should read this document as a proposal from an interested party, and should discount it accordingly.

The establishment table includes a Director / Chief Executive line at USD 96,000. That figure is a planning assumption I wrote, carrying the same health warning as every other band in §5.2 — it is not a benchmarked market rate and should not be defended as one. It represents 4.9% of five-year payroll, 1.33 times the next-highest post and 2.82 times the median post.

I propose that **the Board sets the Director's remuneration itself, against an external benchmark of comparable director-level appointments, and that the post-holder takes no part in that determination.** If I hold the post, I will not be present for the discussion and will accept what is set.

I make the proposal for a practical reason as well as a proper one. Approving Phase 2 commits roughly USD 15.4 m of the USD 17.4 m five-year total — 89% of the spend sits after the month-15 gate. A Board weighing a commitment of that size should not have to wonder whether the person recommending it has an undeclared interest in the answer. Removing my own pay from my own paper removes the easiest available objection to the whole proposal.

The Board may reasonably go further and separate the roles: founding architect and interim Director through Phases 0–1, with the permanent appointment made competitively at gate 2. I would not resist that. It would mean the decision to commit the larger sum and the decision about who leads the institution are taken separately, which is the cleaner sequence.

---

12. Annex — full establishment

Unit	Post	Count	Gross USD each
Executive	Director / Chief Executive	1	96,000
Quality and Evaluation	Head of Quality and Evaluation	1	60,000
Quality and Evaluation	Evaluation officers	2	30,000
Quality and Evaluation	Data analyst	1	28,000
Quality and Evaluation	Observation and audit officers	2	26,000
Ethics and Data Governance	Ethics and governance lead	1	48,000
Ethics and Data Governance	Data protection officer	1	38,000
Ethics and Data Governance	Research ethics coordinator	1	24,000
Curriculum and Pedagogy	Director of Curriculum and Pedagogy	1	72,000
Curriculum and Pedagogy	Instructional designers	4	32,000
Curriculum and Pedagogy	Clinical content leads	6	54,000
Curriculum and Pedagogy	Assessment psychometrician	1	60,000
Curriculum and Pedagogy	Medical editors and translators	3	22,000
Faculty and Delivery	Track leads	5	50,000
Faculty and Delivery	Certified instructors	10	34,000
Faculty and Delivery	Simulation faculty	2	38,000
Faculty and Delivery	Programme manager	1	34,000
Engineering and Platform	Head of engineering	1	68,000
Engineering and Platform	Full-stack developers	5	36,000
Engineering and Platform	ML and evaluation engineers	3	44,000
Engineering and Platform	Data engineers	2	36,000
Engineering and Platform	DevSecOps engineer	1	42,000
Engineering and Platform	QA engineer	1	28,000
Simulation and Clinical Labs	Simulation centre director	1	52,000
Simulation and Clinical Labs	Simulation technicians	3	20,000
Simulation and Clinical Labs	Standardised-patient programme lead	1	24,000
Simulation and Clinical Labs	Clinical skills tutors	2	30,000
Operations and Registry	Registrar and records	2	20,000
Operations and Registry	Finance and procurement	2	26,000
Operations and Registry	Partnerships	1	34,000

Unit	Post	Count	Gross USD each
Operations and Registry	Communications	1	26,000
Operations and Registry	Monitoring and evaluation officer	1	28,000
Operations and Registry	Administration	1	14,000
	Total	**71**	**2,668,000**

Not counted in the establishment: clinical champions (2 per participating hospital, 0.2 FTE sessional), visiting faculty, fellows, the twelve-member patient and public panel, and the external examiner.

Prepared by Dr Neal Aggarwal, August 2026. Figures computed by `financial_model.py`; re-run after any change of assumption. This document was drafted with AI assistance; the assumptions, the judgements and the responsibility are mine.