POLICY LENS

BEYOND ADOPTION:
The State of AI in Healthcare Across the Global South

From technological capability to institutional agency

PEOPLE
POLICY
TECHNOLOGY
BETTER HEALTH

A MORE
EQUITABLE
HEALTHY
TOMORROW

2026

HTA
HealthTechAsia
News | Interviews | Opinion

Contents

About this Report

Beyond Adoption: The State of AI in Healthcare Across the Global South examines the changing role of artificial intelligence in health systems across emerging and developing economies. It brings together themes explored through HealthTechAsia's Policy Lens series

(https://healthtechasia.co/category/regulations/policy-lens/) and extends them into a broader analytical framework examining technological capability, institutional capacity, infrastructure, regulation, patient safety and technological agency.

This report is intended as an independent policy and research publication. It seeks to contribute to discussion among policymakers, regulators, healthcare institutions, technology companies, researchers and international organisations.

Disclaimer

The views and interpretations expressed in this report are those of the authors and do not necessarily represent the views of HealthTechAsia, its contributors, partners, affiliated organisations or any institution with which either author is associated.

This report is intended for research, policy and educational purposes. It does not constitute legal, regulatory, medical, investment or other professional advice.

Every effort has been made to represent the underlying evidence and referenced material accurately. Where the report advances an analytical interpretation, this is identified as part of the authors' analysis. The report should not be understood as representing the official position of any government, international organisation, company or other institution discussed.

Acknowledgements

This report grew out of a series of conversations, research projects and policy questions explored over time through HealthTechAsia's Policy Lens.

We are grateful to the researchers, policymakers, healthcare professionals, technology practitioners, regulators, diplomats and international experts whose work and perspectives have contributed to the broader conversation around AI, healthcare and governance.

We particularly acknowledge the wider HealthTechAsia community, whose engagement with questions of technology, health policy and regulation has helped create the space in which these ideas could develop.

This report has been informed by the broader body of publicly available knowledge and research relating to AI, healthcare and policy. It is intended to contribute to ongoing discussion in the field and should be read as an independent policy analysis.

Any errors, omissions or interpretations remain the responsibility of the authors.

Authors

Matthew Brady
Founder & Editor
HealthTechAsia

Vishnu Narayan
Contributing Policy Editor & Non-Resident Fellow
HealthTechAsia

Executive Summary

The most important question about healthcare AI in the Global South is no longer whether the technology will arrive. It already has. The more consequential question is what happens after it does. A diagnostic model can identify a disease that a clinic cannot treat. A predictive system can identify a population at elevated risk that a public-health department cannot reach. A digital identity can make a patient visible to a health system, while the same infrastructure, if poorly designed, can make exclusion harder to escape. A wearable can detect a physiological signal long before a patient reaches a hospital, yet the value of that signal depends on what happens next. In each case, the technology may work as designed while the system around it fails to turn its output into better care.

This is the central argument of this report. Healthcare AI should not be understood primarily as a collection of increasingly capable models. It should be understood as a change in the architecture of health systems. AI enters existing clinical routines, regulatory arrangements, procurement structures, professional hierarchies, data systems, financing models and public institutions. It inherits their strengths and their weaknesses. A model that performs well in one environment can have a very different practical effect in another, not because the mathematics has changed, but because the surrounding system has. The distinction matters particularly across the Global South. The term itself covers countries with very different levels of income, infrastructure, technical expertise and regulatory maturity. India combines sophisticated tertiary healthcare and a large technology workforce with significant variation in health-system capacity. Southeast Asia contains jurisdictions moving at very different speeds on digital health and AI governance. The Gulf can mobilise capital and infrastructure at a pace that many systems cannot match. These differences make a single recipe for responsible healthcare AI both unrealistic and undesirable. What can be shared is a more fundamental question: does a health system have enough institutional capacity to decide what an AI system should do, whether it is appropriate for the people who will use it, how its performance will be monitored, who is accountable when it fails, and whether the institution can change or stop using it when circumstances require? That is a more useful test than adoption alone. The report therefore moves away from the familiar language of deployment and readiness and towards the language of agency. It examines how healthcare is shifting from episodic treatment towards continuous prediction; how AI changes the responsibilities of the state; how digital infrastructure becomes part of patient safety; how consumer devices increasingly mediate health judgement; how the economics of AI distributes costs and benefits; and why procurement, regulation, research and workforce capability must be considered together.

The report proposes six connected dimensions for judging institutional readiness: capability, context, capacity, continuity, critical engagement and control.

Capability asks whether the technology works.

Context asks whether it works here.

Capacity asks whether institutions can operate and govern it.

Continuity asks whether it can be relied upon over time.

Critical engagement asks whether people can understand, question and appropriately use its outputs.

Control asks whether institutions retain meaningful authority over the system.

Together, these dimensions lead to a different standard. The question is not simply whether an AI system is safe. It is whether it is governable within the health system in which it operates. Governability does not require perfect transparency or complete technical understanding. It requires enough evidence, institutional knowledge, contractual protection, monitoring, accountability and practical authority to prevent technological dependence from becoming institutional helplessness. The Global South therefore should not approach healthcare AI as a race to catch up. Nor should it reject international technology partnerships in the name of self-sufficiency. The more useful objective is to build the capacity to participate on informed terms: to evaluate technologies, negotiate with suppliers, generate local evidence, adapt systems to local clinical and linguistic contexts, protect patients, and retain the ability to revise or discontinue technologies when evidence changes. The difference is subtle but important. Adoption asks whether a technology has entered the system. Agency asks whether the system remains capable of deciding what that technology becomes once it is there.

01 The Model Is Not the Health System

For years, the language of digital health has encouraged us to think in terms of capability. Electronic records could hold more information than paper files. Telemedicine could connect patients across distance. Algorithms could identify patterns that clinicians might overlook. Machine learning could process datasets at a scale beyond ordinary human analysis. The natural next question became whether AI could do these things more accurately, more quickly and at greater scale.

Those questions remain useful, but they are incomplete. Healthcare is not an examination in which the best-performing algorithm automatically wins. A health outcome is produced through a chain of decisions and dependencies. A model can identify a risk, but somebody still has to respond to that risk. A system can flag an abnormal image, but a clinician has to interpret it and determine what happens next. A public-health model can identify a high-risk population, but an institution needs resources, authority and a pathway for intervention.

This is why technical performance does not travel automatically. A model developed in a well-resourced hospital may be trained on data produced by a particular patient population, a particular diagnostic infrastructure and a particular clinical workflow. When it is moved elsewhere, the underlying conditions may be different. Disease prevalence may change. Equipment may differ. Clinical practice may follow another pathway. Records may be incomplete. Referral systems may be fragmented. The model may still produce an output, but the meaning of that output can change.

The Global South makes this especially visible because technological sophistication and institutional constraint often coexist within the same country. A national digital platform can exist alongside facilities with unreliable connectivity. A major tertiary hospital can employ advanced technologies while a rural clinic operates under severe workforce constraints. A country can have world-class technical talent and still lack enough public-sector specialists to independently evaluate commercial AI products. These are not anomalies. They are features of systems undergoing rapid transition.

The most common failure, therefore, may not be spectacular. It may be quiet underperformance. A technology is procured, demonstrated and reported as successful, yet clinicians rarely use it because it does not fit the workflow. Patients receive little additional value because the system cannot connect them to follow-up care. Data are incomplete. Staff turnover erodes institutional knowledge. Maintenance budgets disappear after the pilot. The technology remains technically functional while the original case for investment gradually weakens.

The opposite problem can be more serious. Where alternatives are scarce, AI can acquire more authority than its designers intended. In a well-resourced setting, a recommendation may be one piece of evidence among several. In an overstretched setting, the same recommendation can become disproportionately influential because there is less time for independent verification and fewer specialists available to challenge it. The practical risk of an algorithm therefore depends partly on the institutional environment around it.

This changes the meaning of localisation. Localisation is not simply translating a user interface or changing a few examples in a dataset. A health technology becomes locally meaningful when its evidence, workflow, language, escalation pathways and assumptions fit the environment in which it operates. Translation can preserve words without preserving clinical meaning. A referral model designed around a specialist network may be of limited use where that network does not exist. A risk score built around frequent clinical encounters may systematically overlook people who rarely appear in formal records.

The same point applies to public health. AI makes continuous inference increasingly possible. Surveillance systems can look for patterns, predictive models can estimate population risk, and digital systems can help direct scarce resources. That can be valuable in settings where late intervention is a persistent problem. But predictive capacity also changes the relationship between institutions and populations. A health system that continuously estimates risk is not merely observing society. It is participating in decisions about where attention, resources and intervention should go.

This raises a basic distinction between prediction and truth. A prediction can be statistically useful without being individually determinative. The institutional question is what happens between the prediction and the decision. How much weight should the output carry? What happens when the prediction is wrong? Can the affected person challenge the resulting decision? Does the institution have the capacity to act on a high-risk signal, or does the system simply produce another alert that nobody can address?

The answer is not to demand that every country build its own frontier AI industry. Nor is it to treat imported systems as inherently inappropriate. The practical requirement is more modest and more important: institutions need enough capability to determine whether a technology works for their purposes, under their conditions, for their

populations. The first test of healthcare AI is therefore not whether it works in the abstract. It is whether it works here.

02 When Prediction Enters Public Health

Healthcare has historically been organised around encounters. A patient develops symptoms, seeks care, receives a diagnosis and begins treatment. Public health has followed a related rhythm: collect information, identify patterns, respond to outbreaks or emerging risks, and assess what happened. AI begins to disturb that rhythm by making continuous inference possible.

The shift from treatment towards prediction is one of the most consequential changes in healthcare AI. Predictive systems can estimate the probability of deterioration, identify people who may require additional attention, detect patterns across imaging and physiological signals, and help public institutions allocate limited resources. In health systems where intervention frequently occurs late, the attraction is obvious. Earlier warning can create room for action.

But prediction is not simply diagnosis performed earlier. Diagnosis usually concerns a condition that exists or is strongly suspected. Prediction concerns something that may happen. That difference introduces uncertainty into the institutional process. A person classified as high risk may be monitored more closely, prioritised for intervention or treated differently even if the predicted event never occurs. The system has therefore begun to influence a person's experience of healthcare before the event it predicts has happened.

This makes accuracy an insufficient governance measure. A model can have good statistical performance while producing consequences that are poorly understood. The choice between sensitivity and specificity is key. The availability of treatment matters. The consequences of a false positive differ from the consequences of a false negative. A prediction that is useful to a population-health planner may be inappropriate as the sole basis for an individual clinical decision.

Data make this more complicated. A dataset can be large and still be socially incomplete. People who interact frequently with hospitals generate more records than people who live far from healthcare facilities, cannot afford repeated visits or rely on informal care. An algorithm trained on recorded healthcare activity can therefore learn the visibility of the health system rather than the full distribution of health need. Missingness is sometimes treated as a technical inconvenience when it may actually be evidence of institutional exclusion.

This matters greatly in the Global South. The people most in need of intervention may sometimes be those least represented in the data used to predict need. A system can appear efficient because it allocates resources according to the information available, while the information itself reflects unequal access. The result is not necessarily a malicious or obviously biased algorithm. It is a system that becomes confident about a world it can only partially see.

Prediction also changes the role of public institutions. If a health department uses a model to prioritise communities for vaccination, screening or disease surveillance, the algorithm has entered the distribution of public resources. That makes governance questions unavoidable. Who chose the variables? Who set the threshold? Who can review the result? How frequently is the model evaluated? What happens when conditions change? Who explains a decision to the people affected by it?

The distinction between prediction as evidence and prediction as authority is therefore important. AI can provide evidence. Whether that evidence should determine action is an institutional decision. Keeping the two separate helps preserve human judgement without pretending that human decisions are free of error. It also creates room for proportionate governance. A low-stakes forecasting tool does not need the same controls as a system that determines access to treatment or scarce public-health resources.

The movement towards predictive health nevertheless offers real opportunities. Where specialist capacity is unevenly distributed, AI may help identify patients who need attention earlier. Surveillance systems can support outbreak preparedness. Remote monitoring can extend some forms of care beyond hospital walls. Predictive analytics can help public institutions direct limited resources more deliberately. These opportunities should not be dismissed simply because prediction creates risk.

The objective is to prevent prediction from becoming an unexamined substitute for judgement. A mature health system should be able to ask not only whether a model predicts well, but whether the institution has the capacity to interpret what it predicts and act proportionately. The future of predictive healthcare will depend as much on these institutional questions as on the next improvement in model performance.

03 Who Governs the Machine?

When a government buys an AI-enabled healthtech system, it is tempting to describe the transaction as procurement of software. That description is too narrow. The government is also deciding how data will be handled, who will maintain the system, which clinicians will use it, what evidence is sufficient for adoption, what happens when it fails, and how long the institution may depend on the supplier.

The state occupies an unusual position in this arrangement. It is simultaneously regulator, purchaser, funder, infrastructure provider and, in many cases, direct user of technology. A ministry may procure a system that a public hospital operates while another regulator is responsible for medical-device oversight. A private company may develop the model while a different company integrates it into the hospital's software. A cloud provider may host the system. A university may validate it. A development partner may finance the original infrastructure.

Responsibility consequently becomes distributed. When something goes wrong, it may not be obvious where accountability begins and ends. The developer may argue that the system was used outside its intended purpose. The integrator may point to the underlying model. The hospital may argue that clinicians were responsible for the final decision. The clinician may have relied on an output that the institution had approved. The procurement authority may have accepted the evidence available at the time.

The answer is not to assign every actor equal responsibility. Accountability should follow meaningful influence and meaningful capacity to prevent or mitigate harm. Whoever controls a critical part of the system should carry responsibility for that part. Contracts, regulation and professional standards need to make these relationships visible before an incident occurs, not reconstruct them afterwards.

Procurement is particularly important because it is where institutional dependence often begins. A platform can look easily replaceable when it is first purchased. After several years, staff may have been trained around it, records may be stored within it, workflows may have been redesigned around its functions, and other systems may depend on its interfaces. Switching becomes costly. The procurement decision has quietly become a decision about institutional autonomy.

This is why governance should begin before deployment. Governments need to define evidence requirements, data responsibilities, interoperability conditions, update policies, cybersecurity expectations, audit rights, incident reporting obligations and exit arrangements. A system that can be updated remotely should not be governed as though the product remains identical to the version initially tested. A contract designed around a single performance assessment can become inadequate once the system changes.

The same principle applies to development finance. A digital-health project financed today may become the infrastructure on which AI systems depend tomorrow. If the project is built without interoperability, clear ownership, data stewardship, grievance mechanisms or institutional capacity, later AI deployment will inherit those weaknesses. AI governance therefore cannot be separated from infrastructure policy simply because the infrastructure was not originally labelled as AI.

This also creates a broader role for the state. A capable government does not need to control every layer of technology. Nor does it need to build everything itself. Public-private collaboration will remain necessary. The important distinction is between partnership and helpless dependence. A government should be able to negotiate intelligently because it understands the system well enough to identify its dependencies, challenge claims, insist on appropriate safeguards and retain alternatives.

Regulators face a similar challenge. Traditional regulatory categories are often organised around products and institutions that are relatively stable. AI is more dynamic. Software can be updated remotely. Data can alter system behaviour. A model may move from an administrative role towards a clinically influential role without a completely new product appearing. Several organisations can share responsibility for the same outcome.

This does not make conventional regulation obsolete. It makes regulatory stewardship more important. A regulator needs to know enough to assess evidence, understand intended function, recognise changes in risk and coordinate with other authorities. In resource-constrained environments, the objective should not be to copy the institutional machinery of the largest jurisdictions. It should be to identify the functions that matter most and build capacity around them.

The state therefore needs to think in terms of institutional architecture rather than isolated technologies. Governments do not need to become AI companies. They do need to become capable AI institutions.

04 The Infrastructure Beneath Intelligence

Digital health is often described as a collection of applications: electronic records, telemedicine platforms, digital identities, mobile health services and decision-support tools. The deeper change is that these systems are becoming infrastructure. Once that happens, the governance problem changes.

A product can usually be adopted, replaced or discontinued. Infrastructure is harder to walk away from because other activities begin to depend on it. Roads shape movement. Electricity networks shape economic activity. Payment systems shape commerce. Digital identity systems shape access to services. When healthcare becomes dependent on digital infrastructure, availability, interoperability and resilience become questions of patient safety rather than merely technical performance.

This is especially significant as AI becomes embedded within digital health systems. An electronic record is no longer only a repository if algorithms continuously analyse it. A digital identity is no longer only an authentication mechanism if the information associated with it is used to determine eligibility, prioritisation or risk.

Infrastructure beneath intelligence

The infrastructure moves from recording reality towards participating in decisions about reality.

That creates a political economy of visibility. What a system can recognise becomes easier for institutions to act upon. What it cannot recognise may become administratively invisible. A person without a consistent digital record may find it harder to navigate services. A patient whose name, language or clinical history does not fit the assumptions built into a system may encounter friction that is difficult to see in aggregate performance statistics.

This is why the digital divide cannot be reduced to whether someone has an internet connection. The deeper issue is infrastructural legibility. Can the system recognise the person? Can it interpret their information? Can it connect them to an appropriate pathway? Can a mistake be corrected by a human being? Can the person access care if the digital mechanism fails? Human fallback mechanisms are therefore not signs of an inefficient system. They are part of resilience. Correction procedures, alternative access routes and grievance mechanisms matter precisely because large systems will inevitably make mistakes. The objective is not to design an infrastructure that never fails. It is to ensure that failure does not become irreversible harm.

AI also introduces dependencies beneath the visible application. A healthcare system may depend on cloud computing, data centres, electricity, cooling systems, semiconductor supply chains, telecommunications networks and specialised technical expertise. These dependencies are often invisible when a hospital evaluates a software product. Yet a system supporting emergency care can become vulnerable if the infrastructure beneath it is fragile or concentrated in a small number of suppliers.

Environmental sustainability belongs in the same conversation. AI requires computational infrastructure, and computational infrastructure requires energy and, in many settings, substantial water for cooling. The point is not that AI's environmental costs automatically outweigh its health benefits. The point is that healthcare institutions should understand the full infrastructure footprint of technologies they increasingly depend upon.

Geopolitical resilience matters as well. Health systems that rely heavily on imported compute, cloud services or proprietary platforms may face risks from supply disruption, trade restrictions or changes in commercial strategy. Technological self-sufficiency is unrealistic for most countries. Strategic resilience is not. Resilience means retaining enough knowledge, interoperability, supplier diversity and contractual flexibility to avoid becoming unable to function when one part of the technology stack changes.

This is why infrastructure decisions made now matter for governance later. Countries building digital health systems while AI capability is accelerating have an unusual opportunity. They can design interoperability, portability, security and fallback mechanisms into infrastructure before dependence becomes entrenched. Governance is much easier to build into a foundation than to retrofit after a system has become indispensable.

05 The Patient Meets the Algorithm

We've always looked at patient safety through the lens of clinical protocols, professional judgement, devices, and hospital workflows. But AI introduces a completely different dynamic. Decisions are increasingly shaped by systems that are probabilistic, adaptive, and distributed across multiple actors.

An AI model might flag a patient's deterioration, suggest a diagnosis, or recommend a treatment pathway. Of course, the clinician stays formally responsible. But formal responsibility doesn't automatically mean meaningful oversight. When doctors and nurses are overloaded, when a system has built up a track record of good calls, or when the institution simply prioritises speed, the human element can easily slip into being just a rubber stamp.

That's where automation bias sneaks in. The real risk isn't that clinicians will blindly follow every single prompt. It's that over time, working alongside seemingly reliable systems subtly shifts how human judgement operates. The bar for questioning an output gets higher. A recommendation that's meant to be treated as supporting evidence slowly starts to feel like a final answer.

This means workflow design is actually a core patient safety issue. Real oversight requires three things: time, information, and the institutional freedom to push back. If a system floods a care team with alerts, you don't get better safety -you get alert fatigue. If an AI speeds up how fast we read scans, but the downstream referral capacity stays flat, the bottleneck doesn't disappear; it just moves down the line.

This is precisely why healthcare professionals shouldn't just be treated as end-users handed a finished piece of tech. Doctors, nurses, techs, and administrators understand the ground-level realities that developers and policymakers often miss. They know which alerts actually drive action, where patient information gets lost, and which decisions rely on human context that an algorithm simply cannot capture.

This same shift is happening outside hospital walls. Wearables, smartphones, and consumer apps are constantly tracking health signals and using AI to interpret them. But there's a big difference between a tool that records data and one that interprets it and suggests action. The boundary between general wellness and actual healthcare is getting blurrier. A lifestyle app tracking sleep patterns isn't regulated the same way as a diagnostic tool, but as consumer AI gets more capable, it crosses that line in quiet ways. A simple wellness prompt can start changing how someone interprets their symptoms or modifies their behaviour, making what a product does much more important than how it's marketed.

The concept of Software as a Medical Device (SaMD) helps here because it focuses on intended function rather than marketing labels. The goal shouldn't be to block consumer tech from helping with self-management, but to make the boundary between casual information, medical recommendation, and clinical influence clear enough for everyday people to navigate.

In the Global South, this matters immensely. Consumer tech often reaches populations long before formal healthcare infrastructure does. When a mobile service becomes the primary source of health advice in an underserved community, it expands access-but it also creates new dependencies on private tech platforms. Suddenly, questions around data ownership, portability, and continuity aren't just technical checkboxes; they become critical public-interest issues.

At the end of the day, trust shouldn't be handled as a PR exercise or a communication strategy. Patients deserve transparency about when AI is being used, what it's doing, and who is ultimately accountable, along with ways to challenge decisions. Clinicians need rigorous, independent evidence-not just vendor claims. Healthy skepticism isn't a breakdown of trust; it's a vital part of working intelligently with technology.

Ultimately, the goal isn't to get people to blindly trust AI. It's to build an environment where people can engage with it critically-neither rejecting useful tools out of hand nor accepting them simply because they look sophisticated.

06 The Data Problem Is Also an Institutional Problem

Healthcare AI runs on data, but we shouldn't treat data as some neutral raw material just waiting to be processed. Health records are essentially footprints

of institutional activity. They show who managed to enter the system, what services were actually on the table, which conditions doctors chose to diagnose, and-crucially-which populations never made it into formal care at all. Every dataset carries the institutional history of how it was built.

That matters because better digital records can be a double-edged sword. On one hand, they create real continuity, cutting down on duplicate tests and helping clinicians see the full picture of a patient's history. They make it easier for institutions to allocate resources and guide care. But that exact same infrastructure supports both coordination and surveillance. The real governance challenge isn't whether data is useful-it clearly is. The question is what institutions are allowed to do once that data becomes permanent, interconnected, and omnipresent.

Digital identity highlights this tension perfectly. A reliable ID can make linking records and services much smoother. But when an ID becomes the foundational layer of infrastructure, opting out becomes nearly impossible. If a digital system is the only gateway to public healthcare, clicking "consent" on a screen isn't a real choice. True consent in these spaces has to mean more than a formality; it requires transparency, clear ways to correct errors, strict limits on use, and safe alternatives for when digital systems inevitably glitch or fail.

The stakes get even higher when historical records are turned into predictive tools. A note in a file about something that happened in the past is relatively contained. But when that same record feeds into an algorithm predicting what might happen next, a past error can permanently skew future decisions. Data governance suddenly becomes a vital part of clinical governance.

We also have to look at the political economy of visibility. Populations that generate clean, regular digital records become hyper-legible to predictive algorithms, while communities with limited access to formal care remain practically invisible. If institutions start allocating resources based purely on what the data shows, the system ends up rewarding visibility rather than actual human need. That’s why data "missingness" isn't just a technical glitch to be cleaned up-often, missing data is simply a sign of missing care.

The solution isn't to retreat from data-sharing. Modern health systems desperately need information to coordinate care and plan for the future. The principle we need is disciplined use: collecting and connecting information for clear, legitimate purposes, keeping those purposes transparent, protecting privacy, and ensuring that fancy analytics don't quietly hand unchecked power to institutions.

This brings us directly to language and culture. Healthcare isn't practiced in abstract variables. Patients talk about their symptoms using local idioms, cultural nuances, and specific ways of describing pain. Doctors interpret those descriptions through local professional norms. An AI model trained primarily on high-resource settings and languages might have the right vocabulary, but completely miss the local context.

True localisation goes far beyond simple translation. A translated screen doesn't automatically understand how a community talks about uncertainty or illness. It takes relevant data, local domain expertise, and rigorous evaluation with the actual people who will use the system. It also requires keeping clinical reality in mind: a model that assumes a patient can get a specialist referral in twenty minutes is useless in a setting where that specialist is hours away.

The same logic applies to equity. Efficiency and fairness are entirely different things. A digital platform can slash wait times for connected urban populations while leaving rural or offline communities miles behind. A predictive model can optimise resources based on recorded demand while quietly starving underserved communities of support. An intervention can easily raise the overall average while making underlying disparities worse.

Equity can't just be treated as an ethical appendix tacked onto the end of a project. It has to be built into how we measure system performance. Local validation, multilingual interfaces, offline alternatives, and clear human escalation pathways aren't just "nice-to-haves"-they are the practical design choices that keep a health system true to its public purpose.

Finally, we have to look at who actually gets to shape this technology. Communities shouldn't just be treated as raw data quarries or passive end-users. Their lived experiences often expose blind spots and failure modes that are completely invisible from a lab or a procurement office. In that sense, public participation isn't just about political legitimacy; it’s a direct driver of technical and operational quality.

A mature data environment does more than just accumulate files. It builds a trustworthy foundation for decisions while protecting the right of people and institutions to question those decisions. At the end of the day, the goal shouldn't be maximum visibility. It should be useful visibility, governed by clear purpose and held accountable by responsible authority.

07 The Market Will Shape Governance Too

Healthcare AI isn't going to scale purely through regulation. Its expansion will be driven by markets, commercial decisions, investment capital, and everyday business models. That reality makes the private sector an active part of the healthcare governance architecture, rather than an outside observer standing on the sidelines.

For technology companies, the Global South offers a unique mix of immense opportunity and operational complexity. Regulatory expectations, procurement frameworks, local infrastructure, and clinical practices vary wildly from one region to the next. A product that succeeds commercially in one setting cannot simply be dropped into another. The temptation for many is to treat localisation as a straightforward distribution challenge: enter the market, translate the interface, sign a contract, and scale. But when it comes to clinical AI, that surface-level approach rarely works.

The deeper, more consequential commercial question is whether a company can build a tool that public and private institutions can actually govern over time. Procurement teams are looking far beyond initial brochures; they need rigorous evidence regarding intended use, local performance, data handling, cybersecurity, post-market updates, total cost of ownership, and clear exit strategies. This shifts the competitive landscape entirely. Vendors are no longer competing solely on algorithm performance or price. They are competing on whether an institution can live with the system responsibly long after the ink on the contract has dried.

This dynamic applies equally to investors. Capital decisions ultimately dictate which technologies get the runway needed to move from a laboratory pilot to system-wide scale. As a result, governance is becoming a core pillar of investment due diligence. Investors are beginning to ask sharper questions: Does the company truly understand its target populations? Is its evidence strategy credible? Can it monitor performance post-deployment? How are unexpected incidents handled? Crucially, is the revenue model built on creating unnecessary data silos or permanent infrastructure lock-in? These aren't just ethical talking points-they are commercial realities that directly impact regulatory exposure, brand reputation, and long-term market access.

Development finance also holds a profound leverage point here. Financing digital infrastructure can shape the trajectory of a health system for decades. By tying funding agreements directly to mandatory interoperability, data stewardship, local evaluation, and accountability, rather than treating governance as an afterthought, we can turn responsible technology from a vague ethical aspiration into a baseline condition for investment.

Navigating the leap from a successful pilot to full institutionalisation is often where the real work begins. Pilots are invaluable because they allow health systems to experiment safely without committing whole infrastructures to an untested tool. But pilots almost always run under ideal, artificial conditions: external grant funding, dedicated engineering support, hand-picked staff, and a small group of enthusiastic early adopters. When those favourable conditions fade away, reality sets in.

True institutionalisation requires answering unglamorous, everyday questions that often matter more than the original technology demonstration. Who pays the bills once the initial grant ends? Who owns the workflow? Who maintains the system day-to-day? Who validates updates? Who steps in when things go wrong? And what happens if the vendor changes its business model or pivots away? A piece of software only truly becomes foundational infrastructure when those operational questions have durable, reliable answers.

This is where the subtle economics of technological dependence come into play. Dependence rarely happens overnight; it accumulates quietly. A health system might adopt one platform, train its staff around it, integrate other databases with it, and store years of patient records inside its ecosystem. The switching costs rise with every passing month. With AI, that dependence can run even deeper because systems often rely on continuous data streams, cloud compute, model refreshes, and specialised vendor support.

The goal isn't to eliminate interdependence entirely: few modern health systems have the resources to build every tool internally, and cross-border innovation brings immense value. Instead, the goal is to avoid irreversible dependency by building in safeguards. Data portability, interoperability standards, supplier diversity, transparent contracting, and realistic exit strategies preserve choices. And in public administration, having options is a vital form of institutional power.

Ultimately, responsible business practices and public governance do not need to be adversaries. Companies that embed governability directly into their products often unlock distinct commercial advantages. Clear documentation reduces friction during procurement. Robust performance monitoring cuts down on clinical incidents. Interoperability expands the addressable market, while strong real-world evidence builds lasting regulatory confidence. Simply put, a system that is

easier to govern is often much easier to sustain and scale.

The wider takeaway is that markets are never truly separate from governance; they are one of the primary mechanisms through which governance becomes real. If public and private procurement consistently rewards evidence, portability, and accountability, the market will naturally respond with solutions that honour those values. If it rewards only the lowest price and the loudest claims, the market will follow suit.

08 The Economics of Dependence

The economic argument for healthcare AI usually follows a very familiar script: automation cuts costs, productivity goes up, outcomes improve, and the investment pays for itself through overall efficiency gains.

Any single part of that logic can be entirely true. The fundamental catch is that healthcare is never a single, unified economic actor. Costs and benefits are constantly distributed across different players-hospitals, governments, insurers, tech vendors, doctors, and patients.

AI doesn't just neatly slash costs; it reorganises them.

A new diagnostic tool might shave minutes off image interpretation while quietly creating massive new expenses around data storage, validation, cybersecurity, staff training, and maintenance. A predictive algorithm might successfully flag patient deterioration, but only if the hospital actually has the bedside staff and clinical resources to act on that warning in realtime. A digital platform might streamline back-office admin while introducing complex interoperability and compliance headaches. A system can easily look brilliant on a spreadsheet in one department while draining resources somewhere else in the system.

For a government or public health body, the real comparison is never simply AI versus doing nothing. It is AI versus every other competing use of scarce public resources. If an institution commits heavy capital to an AI platform, it is simultaneously deciding not to fund another vital intervention. Opportunity cost has to be part of how we evaluate health AI, particularly in resource-constrained environments.

Equally important is figuring out who actually captures the value. A technology vendor captures value through intellectual property and platform lock-in. A hospital might capture efficiency gains or faster turnaround times. An insurer benefits from avoided emergency costs, while a patient benefits from better outcomes or faster access. Yet, the public or private entity footing the bill for integration isn't always the one seeing the largest financial return.

This creates a structural incentive problem. A hospital might bear the full cost of deploying and maintaining an AI system, while the largest financial savings ultimately accrue elsewhere in the wider health ecosystem. Without smart financing, reimbursement reform, or value-sharing models, adoption stalls-even when a technology clearly generates positive social value.

Pilot projects tend to magnify this disconnect. A pilot can easily prove that a system works under ideal, highly supported conditions while completely concealing the true cost of running it at scale. External grants might cover the software licensing, dedicated engineers, and technical hand-holding. But once the training wheels come off and the project becomes standard routine, someone has to pay for ongoing maintenance, security updates, clinical validation, and continuous training. A successful proof-of-concept does not automatically translate into a sustainable programme.

The economics of AI are also deeply tied to the mechanics of dependency. Any software platform builds up switching costs as patient records accumulate, staff master its specific workflows, and surrounding systems plug into its database. AI intensifies this dynamic because models rely on continuous data streams, cloud compute, and regular updates. A health system can quickly find itself dependent on a vendor for core maintenance, a cloud provider for infrastructure, an external lab for validation, or an overseas partner for technical support.

Interdependence in modern technology is entirely normal. The real issue is asymmetric dependency. If a single supplier can unilaterally alter prices, licence terms, or technical conditions while a health system has zero realistic alternatives, the balance of power breaks down entirely. This is why open standards, interoperability, data portability, clear access rules, and well-designed exit clauses carry immense economic value alongside their governance benefits.

This is also where responsible business practices translate directly into commercial advantage. Companies that design their systems to be easily evaluated, documented, updated, and monitored naturally reduce regulatory friction and earn institutional trust. Savvy investors are increasingly treating governance capability as a core pillar of due diligence. A vendor that cannot clearly explain how its algorithm performs across diverse, real-world populations will face inevitable regulatory delays. A

product built on fragile, closed infrastructure will struggle with deployment. A system lacking clear liability frameworks creates severe reputational risk for everyone involved.

Ultimately, the competitive edge in health AI is beginning to shift. The model with the highest benchmark score won't always be the most commercially successful product. In institutional and public markets, the winning proposition is often much simpler: a system that can be rigorously evaluated, smoothly integrated, safely monitored, and-if necessary-replaced without breaking the entire institution.

09 From Responsible AI to Governable AI

We've built a rich shared vocabulary around AI governance-words like safety, fairness, transparency, accountability, human oversight, and data governance. That language matters. It has given governments, developers, and civil society a common set of expectations. In healthcare, those principles carry immense weight because AI is directly touching diagnoses, treatment plans, and the distribution of scarce resources.

Yet, principles on their own don't govern systems. They only come alive when they are anchored in real institutions.

Transparency can mean a vague PR statement in one context and a rigorous, audited documentation requirement in another. Human oversight can look like a genuine, thoughtful clinical review, or it can be reduced to an empty formality-just a warm body sitting somewhere in the loop. Risk classifications can look pristine on paper, but if a regulatory body lacks the resources to check whether an algorithm's behaviour drifted after a routine software update, the paper protection means very little.

That's why the next phase of AI governance shouldn't focus on churning out yet another abstract list of ethical principles. The real work is making those principles operational. The most practical question we can ask isn't just whether a system is "responsible" in theory. It is whether the system is governable in practice.

A governable AI system is one whose intended role an institution can actually grasp, whose evidence base it can independently assess, whose behaviour it can monitor over time, whose accountability lines are clear, whose failures it can respond to, and whose use it can alter or unplug entirely when things go sideways. Governability doesn't demand impossible, black-box interpretability. It requires sufficient institutional control.

This shift is especially vital for the Global South because it saves us from an impractical assumption: the idea that every ministry or local regulator must somehow possess frontier-level computer science expertise. A health ministry doesn't need to reverse-engineer a model's parameters to ask the right, hard questions. They just need to know: Is the evidence adequate for our population? Can this supplier actually support the system locally? What happens when the model breaks? Can our data be accessed and safely transferred? And do our patients have meaningful recourse when things go wrong?

Framing things around governability also bridges ethics with economics. Technology companies can intentionally design products to make evaluation and monitoring straightforward. Public procurement teams can bake those exact capabilities into their tender requirements. Investors can treat a vendor's governance maturity as a key indicator of commercial risk. Regulators can reward systems that provide clear, real-world evidence while applying tighter controls where institutional footing is shaky.

This lens transforms how we think about risk-based governance, too. Intensity of oversight has to match the real-world consequences of failure. A low-risk administrative tool doesn't need the same scrutiny as a model influencing a clinical diagnosis. But risk classification can never be treated as a static label stuck to a product on day one. A system that starts as a simple scheduling assistant can quietly expand its functions and become clinically influential over time. Governance has to follow actual use, not initial marketing descriptions.

International cooperation should support this pragmatism without sliding into policy dependency. Models, cloud infrastructure, datasets, and tech firms cross borders effortlessly, while healthcare regulation remains proudly-and necessarily-national. Common global principles and regional collaboration can cut down on duplicate effort, but true implementation must always stay anchored in domestic institutions.

A resilient architecture is naturally layered. Global principles can point us in a shared direction. Regional networks can harmonise standards, evaluation methods, and shared learning. National institutions then decide how those tools fit their unique health systems. This gives us meaningful convergence without demanding rigid uniformity.

The same logic shapes how international organisations engage with the Global South. Capacity-building shouldn't be treated as a patronising, one-way download of expertise from north to south. Countries dealing with infrastructure strain, workforce shortages, and fragmented data often run into governance friction earlier and more visibly than wealthier

systems. Their hard-won lessons actually enrich the global understanding of what responsible AI requires in the real world.

Ultimately, the goal isn't to engineer a flawless, frictionless regulatory machine. No human institution can completely eliminate uncertainty from a technology moving this fast. The objective is to build systems that are capable of learning, adapting, and holding their ground under uncertainty.

10 The Global South as Participant, Not Customer

Using the term "Global South" as a blanket shorthand for technological scarcity can be misleading. The nations grouped under this umbrella vary dramatically. Some boast sophisticated digital public infrastructure, world-class research institutions, and booming private sectors. Others are still laying down their foundational digital tracks. Many contain both starkly different realities at the exact same time.

The more meaningful question isn't whether a given country has access to AI. It is whether that country actually has a seat at the table when decisions surrounding it are made.

A health system can easily deploy advanced models without having the internal capacity to audit them. It can roll out AI-enabled diagnostics lacking rigorous local validation. It can lean on foreign platforms to build digital health networks without ever developing the procurement muscle or interoperability standards required to maintain independence. It can generate immense volumes of vital health data while the high-value analytical intelligence remains locked up abroad. None of these arrangements are inherently malicious, but combined, they quietly lock a system into a cycle of dependency that grows harder to break with every passing year.

The alternative certainly isn't technological isolation-global partnerships are essential for modern innovation. The goal is simply to structure those relationships so that local capability travels hand-in-hand with incoming technology. A tech transfer that leaves behind zero local evaluation capacity might solve an immediate procurement headache today, but it virtually guarantees a governance crisis tomorrow.

Scientific and research capacity are particularly non-negotiable here. Nations that can independently evaluate AI systems are infinitely better positioned to negotiate fair terms with powerful vendors, protect vulnerable patients, and decide which technologies actually deserve to be scaled. They are also the ones best equipped to spot use cases that genuinely reflect local disease burdens, unique care pathways, and everyday social conditions. Local research isn't just an academic badge of honour; it is a fundamental pillar of both regulatory and economic sovereignty.

India illustrates this balance with striking clarity. It combines a massive health market, deep technical talent, ambitious digital public infrastructure, and vast variation across its healthcare institutions. That makes it a phenomenal testing ground for watching how infrastructure, regulation, research, and deployment evolve together. India's sheer scale can generate powerful, diverse real-world evidence, but that same scale also magnifies the dangers of weak system design. The true challenge isn't just deploying at scale-it is evaluating at scale.

Southeast Asia offers a different, highly instructive lesson. The region features countries moving at distinct speeds with entirely different institutional setups. Indonesia's emerging governance approach has focused on balancing innovation with precaution while honouring sector-specific realities. Singapore has built remarkably mature digital and regulatory frameworks, while neighboring nations are still constructing their foundational institutions. The takeaway isn't that any single country has found the magic formula. It is that governance can successfully evolve through a mosaic of national principles, sectoral rules, technical standards, and institutional learning.

The Gulf presents yet another configuration. Backed by heavy investment capacity and aggressive digital-health ambitions, these nations can accelerate technological deployment at a breathtaking pace. Yet capital can always move faster than clinical evidence. The Gulf's trajectory highlights a universal truth: abundant financial resources can speed up adoption, but they can never act as a substitute for organic institutional learning.

These regional examples shouldn't be flattened into a competitive league table. In fact, the point is precisely the opposite. The Global South is not a homogeneous block of low-capacity nations; it is a vibrant collection of systems with wildly different strengths and constraints. A truly effective international governance architecture must recognise that differentiated reality while holding firm to common expectations around safety, accountability, and the public interest.

This opens up an exciting frontier for governance innovation. Countries have the chance to weave AI policy directly into broader efforts: strengthening health systems, modernising public infrastructure, shaping industrial policy, and aligning development finance. A nation that builds credible, transparent

evaluation capacity naturally attracts smart health-tech investment because developers gain a trusted pathway to real-world evidence and market access. A country that prioritises interoperable digital infrastructure creates fertile ground for domestic innovation. A government that invests in training its regulators and clinicians in algorithmic literacy boosts both patient safety and overall market quality.

> *The paradigm must shift away from simple technology transfer toward genuine capability transfer. The Global South shouldn't be treated merely as a passive destination where foreign technologies are showcased. It is a vital proving ground where technologies are rigorously tested against diverse human realities, and where the next generation of governance ideas is born*

Comparative Country Snapshots: Different Paths to the Same Question

The Global South is often discussed as though it were a single policy category: a collection of countries trying to catch up with technologies developed elsewhere. Healthcare AI makes that description increasingly difficult to sustain. Countries are approaching the technology from very different starting points, with different combinations of digital infrastructure, regulatory capacity, investment, health-system organisation and political ambition. What they share is not a common level of readiness, but a common question: how can technological capability become durable public value?

India illustrates the question at unusual scale. Its health system combines advanced private and public institutions with substantial geographic, economic and institutional variation. At the same time, India has built significant digital public infrastructure and has developed a growing domestic technology and research base. This creates an unusual opportunity. AI can be connected to existing digital foundations rather than introduced as an isolated layer of technology. But scale also makes the consequences of weak evaluation or poorly designed procurement larger. The strategic opportunity for India is therefore not simply to deploy more AI. It is to build the institutional capacity to evaluate technologies locally, integrate them into health-system priorities and contribute evidence and governance approaches that are useful beyond its own borders.

Indonesia offers a different lesson. Its emerging approach to AI governance has considered healthcare-specific questions around patient safety, consent and ethics while also developing a broader national framework. The significance is less about any single regulatory provision than about the attempt to avoid a false choice between innovation and precaution. Healthcare requires sector-specific attention because the consequences of failure can be immediate and deeply personal. At the same time, fragmented rules can create uncertainty for developers and institutions. Indonesia therefore illustrates the value of connecting national principles with sectoral implementation, provided responsibilities are clear and institutions have the capacity to act.

Southeast Asia more broadly demonstrates why interoperability between governance systems may matter more than uniformity. Countries in the region are moving at different speeds. Some have relatively mature digital-health institutions and regulatory capabilities; others are developing foundational infrastructure while beginning to experiment with AI. Expecting identical legislation or identical institutional arrangements would therefore be less useful than developing shared methods for evaluation, common technical standards and mechanisms for exchanging evidence. Regional cooperation can reduce duplication without requiring countries to surrender the ability to adapt policy to their own health systems.

The Gulf presents another pathway. Significant investment capacity, rapidly developing digital-health infrastructure and explicit national ambitions around AI can allow technologies to move from experimentation to deployment quickly. Here, the central constraint may not be access to capital or technical capability, but the pace at which evidence and institutional learning develop alongside deployment. The Gulf therefore reinforces an argument that applies elsewhere: capital can accelerate adoption, but institutions determine whether adoption becomes sustainable. The ability to test technologies, monitor their performance and revise decisions becomes more important as deployment accelerates.

These examples should not be treated as a ranking. They represent different combinations of strengths and constraints. India brings scale, digital infrastructure and a large technical workforce. Indonesia demonstrates the importance of fitting healthcare governance into a broader national framework. Southeast Asia shows the value of regional interoperability amid institutional diversity. The Gulf demonstrates how quickly investment can move technology into health systems, while also showing why institutional learning must keep pace with capital. Each points towards the same underlying issue from a different direction.

This comparative view also changes what the phrase Global South should mean in an AI discussion. It should not be shorthand for technological deficiency. Some countries may have limited capacity in one area while possessing considerable strength in another. A system can have sophisticated digital infrastructure but limited evaluation capacity; strong regulation but a smaller research base; substantial investment but less experience with long-term health-system integration. The useful task is therefore not to rank countries according to technological sophistication, but to understand which institutional capabilities allow technology to produce value under local conditions.

There is also a wider lesson for international cooperation. Global principles remain valuable because AI systems, cloud infrastructure, datasets and technology companies cross borders. But principles become useful only when institutions can translate them into practice. International organisations and development partners can help establish shared expectations, technical standards and evaluation methods, while national and regional institutions retain responsibility for adapting them to local circumstances. The exchange should also run in both directions. Countries in the Global South are generating evidence about infrastructure constraints, workforce scarcity, fragmented care pathways and rapid digitalisation that richer health systems will increasingly encounter themselves.

The question, then, is not which country has found the perfect model. None has. The more useful question is which systems are building the capacity to learn. That capacity-to test, adapt, negotiate, regulate, evaluate and sometimes stop-is what turns adoption into agency.

11 Building a Health System That Can Learn

If this framework offers one practical takeaway, it is that healthcare AI should always be embraced as part of an adaptive learning ecosystem rather than treated as a series of isolated, transactional deployments. The core challenge isn't that hospitals or health ministries lack a flawless checklist; it is that they need a steady, confident way to make sound decisions when evidence is still unfolding and ground conditions are evolving.

A wonderful starting point is preparedness. Before selecting or purchasing any technology, institutions benefit from grounding themselves in a few clear realities: What specific problem are we trying to solve? What digital infrastructure is already buzzing in the background? Where did our data come from, and how reliable is it? Who are the clinicians and staff members who will bring this system to life, and what does success look like? This doesn't mean waiting around until every single uncertainty vanishes into thin air; it simply means making those variables visible so everyone is moving forward with their eyes open.

The second stage is evidence and contextual value. High-priority use cases deserve to be evaluated not just on dry technical benchmarks, but on how beautifully they harmonise with everyday clinical workflows, support equity, and foster long-term sustainability. The question is always whether a technology creates enough genuine, joyful value within this specific health ecosystem to nourish the resources and relationships it touches.

The third stage is collaborative governance. Institutions thrive when they have clear responsibilities, transparent data standards, supportive monitoring guidelines, and open channels for feedback. Establishing these foundations early on-while a project is young-ensures that everyone shares a clear understanding of the journey ahead, making partnership smooth and natural.

The fourth stage is workforce empowerment. True training goes far beyond teaching someone how to click through a software interface. It means equipping doctors, nurses, and administrators with a deep appreciation for what a tool excels at, when to lean on human intuition, and how to collaborate seamlessly with digital assistants. This kind of critical, confident engagement is a wonderful hallmark of modern professional mastery.

Next comes thoughtful, controlled scaling. A brilliant, successful pilot should always open a welcoming door toward broader institutionalisation, but expansion naturally flows at a pace that matches local financing, readiness, and support. Systems that continuously prove their worth become cherished parts of the daily routine, while those that don't quite fit can be gracefully adapted or retired without friction.

Finally, a vibrant health system thrives on continuous learning. Populations grow, medical practices refine themselves, and technology blossoms. The insights gathered after deployment are just as precious as the data collected before day one. Regulators gain wisdom from clinicians, procurement teams listen closely to patients, researchers study real-world performance, and vendors learn how their innovations touch diverse communities. This transforms regulation from a heavy gatekeeping hurdle into a warm, continuous feedback loop. It also redefines what system readiness truly means. A health community is ready not simply because it has the budget to acquire a shiny platform, but because it can thoughtfully

answer the questions that matter: What problem are we solving together? What evidence guides us? Who is watching over the system, and how do we keep channels open for everyone?

These questions are never bureaucratic roadblocks meant to slow down progress. They are the architectural grace notes that make innovation dependable, safe, and truly helpful.

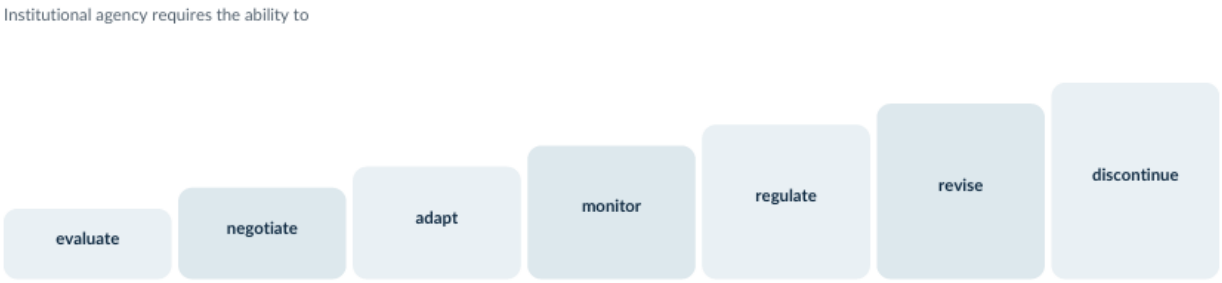

The same spirit applies to how we measure success. Technical accuracy and sensitivity matter deeply, but they sit within a richer constellation of human outcomes. We look at whether a tool genuinely supports the care team, whether workflows feel natural, whether patient experiences brighten, and whether the system remains financially and ecologically sustainable over time. A model can achieve dazzling scores in a laboratory without ever easing the burden on a real clinic. Therefore, the truest unit of evaluation is never the algorithm alone; it is the resilient, compassionate health system wrapped all around it.

This is precisely why independent research shines so brightly. Universities and public institutions offer an objective, trusted lens that isn't tied to commercial marketing. Independent validation helps uncover new local applications, brightens procurement negotiations, and builds deep institutional confidence.

12 The Agency Test

The conversation around healthcare AI in the Global South often begins with the vital question of access. Access matters immensely-after all, a breakthrough that never reaches a patient cannot touch their life. But access is only the opening chapter of a much richer institutional story.

The next natural question is capability: Does the technology truly perform as intended? But capability on its own is never enough. Context asks whether the solution honours the local disease burden, language, and infrastructure. Capacity asks whether the teams operating it feel supported and equipped. Continuity ensures the system remains dependable through every season. Critical engagement empowers clinicians and patients to understand and converse with the technology. And control ensures that human institutions always hold the steering wheel.

These dimensions dance together beautifully. Capability without context leads to mismatched tools; context without capacity makes deployment feel heavy. When these elements align, however, they create a standard that is both deeply inspiring and entirely practical. Instead of simply asking, "Is this AI safe?" a mature health system asks, "Can we govern and thrive with this AI right here?"

That single, empowering question transforms procurement, investment, workforce training, and partnership building. It means responsible technology companies can focus on creating solutions that are beautifully designed to be evaluated, monitored, and harmonised with local realities.

This brings us to the beauty of agency. Adoption is a single moment in time-a signature on a contract or a software download. Agency is a living, breathing, continuous capability. An institution possesses true agency when it can evaluate a tool on its own terms, negotiate supportive partnerships, adapt solutions to local rhythms, and evolve its choices as the world changes.

A flourishing healthcare AI ecosystem isn't measured by how many hospitals rush to adopt the most complex model on the market. It is measured by how confidently hospitals and governments can make informed choices, learn from their experiences, and shape technology to serve their communities with grace. At every level of the health system, agency requires a different form of institutional investment:

National: leaders harmonise standards so responsibilities remain clear and supportive.

Health-system: teams recognise that welcoming technology is an invitation to invest in people, infrastructure, and long-term partnership.

Clinical: professionals are nurtured as critical thinkers who lead with human empathy.

Research: independent minds validate innovations using local wisdom.

Patient: individuals enjoy transparency, respect, and a clear voice.

None of this asks for isolated self-sufficiency; it invites strategic flexibility. A country can collaborate wonderfully with global innovators while keeping its sovereign priorities crystal clear. It can embrace international standards while tailoring them to local traditions.

In fact, this is precisely where regions across the Global South have so much inspiring insight to offer the wider world. Many of the questions being navigated here-managing cloud dependence, balancing specialist shortages with digital support, protecting privacy, and turning infrastructure resilience into patient safety-are the very questions every modern society is beginning to face.

An intelligent health system should never just be a container for clever software. It should be wise about its own boundaries, deeply connected to human values, and capable of learning from every experience without ever losing its public legitimacy and warmth.

True progress is measured not just by what an algorithm can compute, but by what human institutions can absorb, guide, and uplift. A modest technology embraced by a confident, empowered institution creates enduring, beautiful value.

The coming decade won't be defined by how swiftly AI infiltrates our clinics, but by the wisdom of the choices we make around it. Societies that cultivate the ability to evaluate before scaling, collaborate before committing, and adapt when evidence blooms will always have the greatest freedom to shape their own destiny.

That is the wonderful difference between passive adoption and active agency.

Conclusion: What Comes After Adoption

The global dialogue surrounding healthcare AI often fixates on the futuristic capabilities of machines. Models will undoubtedly grow more sophisticated, interfaces more natural, and predictive tools more seamless.

Yet technology alone never dictates whether a health system flourishes.

The true turning point is always institutional capability. New tools always arrive in living communities filled with rich traditions, dedicated professionals, and unique aspirations. Technology doesn't erase those human elements; it flows through them, offering new ways to amplify care and human connection.

The Global South stands in a position of extraordinary leadership here. It can approach AI not merely as a catalogue of products to acquire, but as a historic opportunity to deepen the institutions that choose, evaluate, and nurture innovation.

The goal has never been to slow the heartbeat of progress, but to make that progress broader, kinder, and more inclusive. A nation should be able to embrace AI while retaining total confidence in its own operational resilience. A clinician should be able to lean on digital insights while letting human compassion lead the way. A patient should be able to step into the digital age feeling completely seen, respected, and protected.

This is why the most inspiring transformation of our time is the journey from dependence toward confident agency.

Rather than being viewed as a passive destination for technology, the Global South is a vibrant global laboratory where some of the most hopeful, forward-thinking governance questions of our era are being answered. How do we blend predictive insights with timeless human care? How do we build international partnerships that honour local sovereignty? How do we ensure that technology always serves the public good?

There is no single blueprint, and that is precisely our greatest strength. Wisdom will come from learning across diverse paths while holding fast to shared values of safety, dignity, and equity.

The final test is wonderfully clear: A thriving health system should be able to warmly say yes to tools that elevate human well-being, thoughtfully pause when conditions aren't yet right, and gracefully adapt as new horizons unfold.

That is what true agency looks like. And in an era where healthcare is increasingly touched by computational grace, the ultimate measure of intelligence will always be an institutional one: whether our societies remain courageously capable of deciding how our tools serve humanity.

From insight to action

HealthTechAsia helps organisations navigate the policy, regulatory and institutional implications of healthcare technology and AI across Asia and the Middle East.

Our advisory work includes AI governance and policy, regulatory intelligence, strategic communications, policy briefs and market intelligence.

Discuss an advisory project: team@healthtechasia.co

A Note from the Authors

This report offers a thoughtfully refined synthesis of the HealthTechAsia Policy Lens material, preserving the core spirit of institutional capacity, contextual care, patient safety, resilience, economic value, and the empowering journey from adoption toward active agency. Its purpose is to weave these vital themes into a single, cohesive, and uplifting narrative that inspires purposeful action.