

MAY 2026
ILARIA CALVELLO
SOFIA ARON
NED NAYLOR
DAVID MILESTONE



The Risk of Standing Still: Governing AI in Health Systems Under Pressure

Contents

- 3 Executive Summary
- 6 A Structural Opportunity: How Technology Can Transform Health-Care Capacity
- 9 Rethinking Risk When Adopting Health Technologies
- 12 A Framework for Decision-Making
- 17 Barriers to Technology Adoption
- 20 Recommendations
- 23 Conclusion
- 24 Acknowledgements

Contributors: Vivek Agarwal, Bertrand Assamoi, Mareme Gaye, Patrick Karanja, Magdalene Kariuki, Pule Kumalinga, Ricardo Baptista Leite

Executive Summary

No health system in the world is fully meeting the needs of its population. Irrespective of a country's income, the resources available – both in human and financial terms – are insufficient. The result is that billions of people still lack access to basic health services and, in countries where services are universally available, waiting times are getting longer. The gap is so wide that traditional approaches – incremental productivity gains and training more doctors and nurses – won't close it for decades, if at all.

Recent advances in artificial intelligence offer the promise of a fundamentally different path. If deployed at scale and with the right safeguards, AI could enable a step change in health-system productivity, helping governments to get more out of the finances and clinicians that they have available. Evidence for the clinical effectiveness of AI is substantial and growing, creating a credible opportunity – if current progress continues – for governments to improve access to high-quality health services in years rather than decades, with the potential to save millions of lives in the process.

For this to happen, governments will need to think differently about the risk of deploying technology in health. Given the scale of harm happening in health systems every day, it is imperative that decision-makers weigh up not only the risks associated with introducing new technology, but also those associated with inaction. This can be uncomfortable in health care: "First, do no harm" has been the mantra of doctors for centuries, but when new technology has the potential to provide better care to billions of people, the greater harm may lie in standing still.

This does not mean governments should be reckless. Rather, it means approaching adoption deliberately, assessing both the positive impact and inherent risk in new technologies against current practice. It also requires

putting in place the right safeguards, while recognising that inaction is not risk free. Governments need to implement new regulatory approaches that are proportionate to risk, grounded in health-system realities, able to keep pace with rapid advances in AI and designed to address issues such as bias and inequity. This should be combined with investment in foundational infrastructure and robust data protection, as well as clear engagement with the public and clinicians throughout design and deployment. Governments that are prepared to take a lead will realise the benefits sooner, improving services and delivering better care for their populations.

Recommendations

1. **Reframe regulation around comparative risk.** Explicitly assess the risk and impact of new technologies against the real-world status quo, and set approval thresholds appropriate to the type of technology and its clinical use. Ensure robust, independent oversight of safety, performance and bias. Shift towards adaptive models with strong post-market monitoring.
2. **Set a clear direction and communicate openly.** Be transparent about why AI is being adopted, including both its benefits and the harms already present in current systems. To sustain trust, explain clearly where it will be used, how decisions are made and how risk is managed.
3. **Act decisively on high-impact priorities.** Develop a funded, sequenced plan with clear milestones beyond the pilot stage. Prioritise solutions that tackle core system challenges rather than defaulting to lower-risk, lower-impact tools. Deploy suitable solutions in parallel with infrastructure development to strengthen the wider health system.
4. **Enable clinicians, and co-design with users.** Identify tools that fit clinical workflows and reduce burden, developed with frontline user and patient input. Allocate the necessary training, time and support to build workforce confidence and capability.

5. **Build the foundations and environment for adoption.** Strengthen data quality and security, interoperability and digital infrastructure, while providing clarity on priorities, funding and pathways to market to attract private-sector engagement.

01

A Structural Opportunity: How Technology Can Transform Health-Care Capacity

Across health systems worldwide, waiting lists are growing, inequalities are widening and care remains fragmented. Workforce pressures are intensifying and geography continues to shape health outcomes. The consequences are severe: 8 million people die every year in low- and middle-income countries from conditions that would be treatable with timely, high-quality care,¹ amounting to more than 200 million deaths over the next 25 years if unaddressed.²

The pressures are universal, even if their character differs. In high-income countries, ageing populations, coupled with expanding diagnostics and treatments, are driving demand and costs beyond what resources can match. In the UK about 6 million people are waiting for NHS treatment despite repeated funding increases,^{3,4,5} with health and social-care spending projected to reach levels equivalent to the budgets of the next five major departments combined by 2029.⁶

In low- and middle-income countries, these issues are compounded by more foundational constraints: lower baseline spending, declining foreign aid, limited infrastructure and workforce shortages mean that basic gaps in access and quality remain unaddressed. Kenya, for example, has only 240 radiologists serving a population of more than 56 million, against a potential need of more than 5,000.^{7,8,9} The World Health Organization predicts that by 2030 there will be a shortage of 11 million health-care workers worldwide.¹⁰ No training pipeline or innovative financing approach can close gaps of that scale within a generation.

Doing more of the same will not relieve this pressure on health-care systems. Health-care needs are growing faster than available resources and workforce expansion will take decades. While governments and policymakers consistently recognise health as a major priority, the difficulty

in identifying credible, scalable solutions creates a vicious cycle: limited progress leads to disillusionment, which in turn reduces the appetite for more ambitious reform. Annual productivity improvements of a couple of percentage points matter, but they cannot bridge resource gaps on the scale that health systems are facing.

The success of community health workers shows what is possible when we fundamentally redesign health systems. Their number has grown to more than 6 million globally since their role was prioritised in the 1970s, extending essential care to billions of people and making crucial contributions to improved health outcomes.^{11,12,13,14} This shift required governments and clinicians to make a conscious decision to move beyond conventional approaches that were unlikely to solve persistent access challenges, instead taking managed risks and designing new systems around real constraints.

Today, digital technologies – and artificial intelligence in particular – promise to offer something similar and potentially of even greater impact: the ability to expand the effective reach of scarce clinical expertise, augment capacity and redistribute it more equitably across populations. The impact could be truly transformative: it is conceivable that AI could provide a report on every X-ray taken in a country, enable a consultation with every post-operative patient and provide cognitive behavioural therapy (CBT) to millions of people simultaneously.^{15,16}

The evidence that this is achievable is no longer theoretical. We recently examined AI's potential impact in the context of demographic change in [*Lessons From the UK and Japan: Ageing Populations in the Age of AI*](#), but its applications across health care are far broader and continue to evolve rapidly. In San Diego, an AI algorithm continuously monitoring patients for sepsis – a condition that kills an estimated 11 million people every year – reduced in-hospital mortality by 17 per cent.^{17,18,19} Across 35 countries, an AI-powered chatbot has delivered CBT-based mental-health support to more than 14 million users, with reports of improvements in wellbeing and anxiety symptoms; this demonstrates the potential of AI to expand access to psychological support where services are constrained.²⁰

In these examples, the technology did not replace a functioning system: it reached populations that the existing system could not. The relevant comparison is often not AI against a well-resourced clinical team, but new technology against no care at all, or against a community health worker with minimal diagnostic support. These technologies are emerging precisely where conventional approaches have failed.

At the same time, in some specific domains (including elements of imaging and signal analysis in dermatology and pathology), well-validated AI systems are now matching or exceeding clinician performance.^{21,22,23,24} This shifts the nature of the comparison. Where technology offers greater accuracy, the case for responsible deployment becomes not only pragmatic, but clinically and ethically compelling. These scenarios will only become more common as AI capabilities develop.

02

Rethinking Risk When Adopting Health Technologies

Health governance is shaped by the precautionary principle (where there is uncertainty, action should be avoided where it may cause harm), as well as the commitment to “first, do no harm”. This is foundational and justified. When lives are at stake, caution, robust evidence and disciplined evaluation are essential. However, neither principle requires inaction. The maxim of doing no harm permits treatments with side effects when the alternative is worse: for example, chemotherapy is administered despite its toxicity because untreated cancer is more harmful. In clinical practice, risk is always assessed relative to a counterfactual: the expected outcome without intervention.

In the case of new technologies, governance frameworks often apply a higher burden of proof to emerging tools than to the practices that they are intended to replace. This precautionary bias has been widely observed outside health technology, including in law, risk regulation and technology governance.²⁵ Inaction is often treated as the neutral or safe choice, leaving its risks unquantified and unaddressed.

This double standard is increasingly difficult to justify, given deficiencies in current health systems and the growing evidence for AI effectiveness in specific clinical tasks. The misalignment takes two distinct but related forms.

The first can be understood as a **comparative distortion**: risk assessment focuses on the potential harms of adopting new technologies, while the harm generated by delay or inaction receives less systematic attention. The second is a **proportionality distortion**: governance frameworks apply similar levels of scrutiny across technologies and contexts with fundamentally different risk profiles. Together these dynamics can slow the safe adoption of innovation and misdirect where and how caution is applied.

Comparative Distortion: One-Sided Risk Assessment

Comparative distortion arises when the counterfactual is not systematically incorporated into decision-making. Governance processes rigorously evaluate the risks of adoption but give less structured attention to the risks of delay. As a result, inaction is implicitly treated as safe and proceeds without formal assessment, despite producing predictable harm. This imbalance is reinforced by political incentives: governments are more likely to be held accountable for visible harms associated with their actions than for the frequently unseen harms of maintaining the status quo.

It is patients who often bear the immediate clinical cost of this approach. When new diagnostic or decision-support tools are stalled or not adopted at all, the result is not neutrality but potential harm: delayed diagnoses, missed conditions and preventable disease progression. When emerging technologies can improve the quality and availability of care, delays in adoption carry direct human costs.

Second, delay in AI adoption risks causing the greatest harm to lower-income, more vulnerable and underserved groups. In well-resourced systems, shortages can be partially absorbed through higher clinician density or complementary infrastructure. In under-resourced settings, where expertise is already scarce, these buffers are absent. Technologies that extend diagnostic or decision-making capacity therefore have greater impact in these contexts. Delaying adoption defers that impact, with costs falling most heavily on populations with the least access to quality care.

Delays also affect institutional control. When governments and health systems are slow to adopt and regulate emerging technologies, innovation does not pause; instead it proceeds without oversight. Patients turn to consumer AI tools for medical advice and some clinicians incorporate new tools into workflows without formal governance. Delayed engagement does not prevent adoption; instead it reduces the ability of institutions to shape how it occurs, potentially leading to greater harm.

Proportionality Distortion: Different Risks, Same Scrutiny

Proportionality distortion arises when governance frameworks apply broadly similar evidentiary and regulatory thresholds to technologies with fundamentally different levels of clinical impact, autonomy and failure risk. For example, a medication-dispensing robot, a tool for remote rehabilitation support and an autonomous triage system present very different clinical risks and failure modes, yet governance approaches potentially subject them to comparable levels of scrutiny. The result is a system that delays low-risk, capacity-enhancing applications while waiting for near-perfect safeguards before permitting higher-risk innovations to proceed.

This lack of proportionality has practical consequences. Lower-risk applications – such as administrative automation or appointment scheduling – may be slowed by governance requirements designed for high-stakes clinical decision-making. Oversight capacity can become congested and regulatory attention can be spread thinly across heterogeneous technologies. Rather than strengthening safety, uniform caution can dilute it by failing to concentrate scrutiny where it is most needed.

A proportionate approach would distinguish clearly between levels of clinical impact and exposure to harm, as advocated by a number of international bodies.²⁶ The governance response appropriate for low-impact, low-risk applications is categorically different from that required for high-impact, higher-risk uses such as autonomous diagnostic decision-making. The challenge is not whether to regulate, but how to calibrate scrutiny to both risk and need.

03

A Framework for Decision-Making

The above argument has a practical implication: decisions about adopting new technologies should be based on a structured comparison of their expected impact and risk, relative to current practice. This requires assessment of technology across two dimensions.

1. Impact relative to current practice

The extent to which a particular technology improves outcomes, quality and efficiency relative to what would otherwise occur.

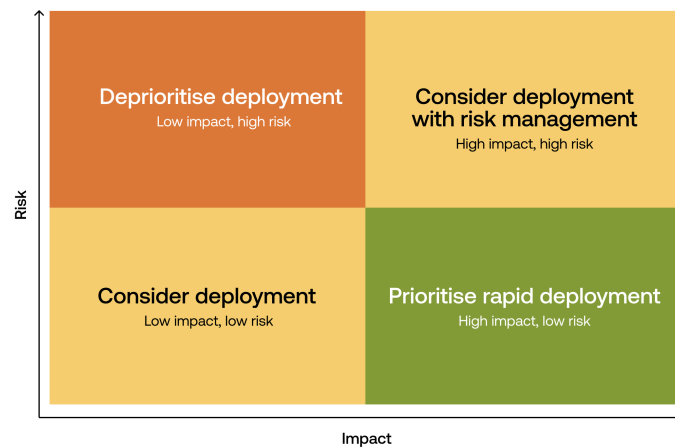
2. Incremental risk of adoption

The additional risks the technology introduces compared to current practice.

The counterfactual – what happens in the absence of adoption – is embedded in both dimensions. It defines both the baseline for impact and the baseline against which additional risk is assessed. Considering these two dimensions facilitates categorisation of technologies and informs subsequent action.

FIGURE 1

Risk and impact framework for technology-adoption decision-making



Source: TBI

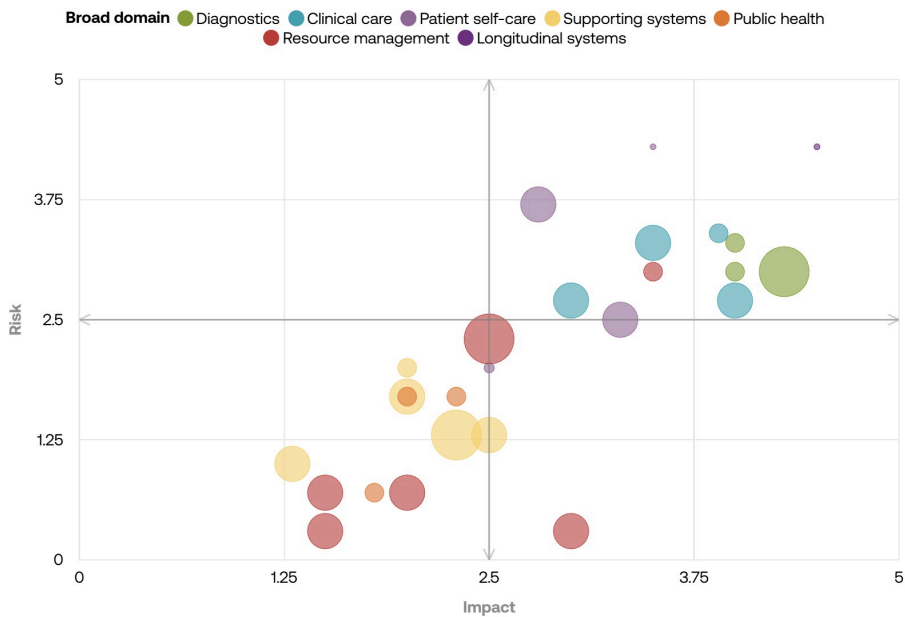
The aim is not to eliminate risk, but to assess it relative to the alternative and act accordingly. This also requires governance frameworks to allow for the fact that some AI tools do not remain fixed after deployment. Systems that evolve with real-world use need greater emphasis on continuous monitoring and post-market surveillance, including periodic reassessment of both their performance and their continued relevance. This ensures that they address the most pressing health-system challenges and remain more effective than the status quo, rather than relying on a single pre-deployment decision.

Practical Implications of the Framework

To illustrate how this framework might be applied in practice we mapped AI use cases according to their expected impact and the incremental risk of adoption, relative to current practice. The aim is not precise quantification but to show how applications span the spectrum of risk and impact.

FIGURE 2

Illustrative assessment: risk and impact of technology use cases across the health system



Source: TBI

Note: This assessment is directional and illustrative, not a precise quantitative-scoring exercise. Informed by available evidence as of March 2026, positioning reflects qualitative judgement of relative risk and impact, while bubble size reflects maturity of the technology. To classify technologies, risk is assessed based on the likelihood of an adverse event occurring, the severity of its consequences and the likelihood of detection before harm results. Impact is assessed based on improvement in clinical outcomes (weighted by the severity of the problem addressed), improvement in productivity, and direct improvement in access to health care (weighted by the population impacted). Scoring is informed by the current evidence on AI effectiveness and safety, and may therefore change as further evidence becomes available. Both dimensions are context-dependent: the same technology may have different risk-impact profiles depending on deployment setting, baseline conditions, availability of alternatives and implementation design. Maturity reflects how ready a technology is for widespread, routine use in health systems, taking into account technical maturity, the strength of the evidence base and implementation feasibility.

Figure 2 reflects this variation. Some technology types, categorised under functional areas using the framework set out in [“Where Do I Start?: How Governments Can Prioritise AI Solutions for Health”](#), combine high impact and relatively low risk, and represent clear opportunities for deployment. Others offer more modest benefits with limited risk and are still worth adopting as

test cases, or to boost productivity with a minimal downside. The most complex cases sit in the high-impact, high-risk category, where the potential to address health-system constraints is greatest but where more measured governance is required.

Our analysis has two important implications for adoption. First, governments should prioritise high-impact applications while aligning adoption with health-system needs. It is often argued that adoption should begin with lower-risk technologies, building confidence before moving to more complex applications. This instinct is understandable and a great deal can be learnt from implementing lower-risk tools. However, our analysis reveals a more difficult reality: many of the technologies with the greatest potential impact also carry greater risk. Given the scale of the capacity gap facing health systems, incremental gains from lower-risk applications are helpful but will not be sufficient to close it. If higher-risk options are delayed too long the result is continued unmet need, which in turn leads to avoidable harm.

Second, adoption should be driven by the greatest need, not the availability of technology. The same application will have a very different impact and risk depending on the baseline with which it is compared, which can vary significantly between and within countries. Decisions should therefore start from an analysis of the specific constraints that a health system faces – diagnostic gaps, priority diseases, limited access to care – rather than from the technologies themselves.

This has particular relevance wherever capacity constraints are greatest. If basic services are limited and workforce shortages are severe, technologies that extend diagnostic or clinical capacity will likely have a greater impact. This does not imply different standards of safety or effectiveness; instead, it reflects the importance of assessing technologies against the realities of current provision and prioritising those that address the most significant constraints.

Contextual fit is essential. However, no tool will be perfectly tailored to every setting before deployment because AI systems, population needs and clinical practice will continue to evolve. In instances where a tool can safely

improve on existing care, the fact that it is not fully tailored to every aspect of the local context from the outset should not automatically prevent adoption. Deployment should be accompanied by monitoring, continued improvement and clear, locally defined safeguards, so that technologies can become better suited to local needs over time.

04

Barriers to Technology Adoption

Even when technologies are effective and justified, implementation often fails because health systems lack the infrastructure to absorb, scale and sustain them. Adoption depends not only on the quality of the technology itself, but also the readiness of the system around it.

FIGURE 3

Key barriers to the adoption of new technologies

Barrier	Why is it important?
Data availability, interoperability and infrastructure	AI systems require large volumes of structured, standardised data. Fragmented records, incomplete digitisation and inconsistent data standards limit functionality and transferability. In many countries, electricity availability, internet accessibility and compute capacity limit the range of AI tools that can be deployed.
Fiscal and procurement constraints	Adoption can require high upfront capital costs with returns realised over longer horizons; complex financing structures (such as blended finance) require technical capability. Traditional procurement approaches may also favour incumbents, prioritise lowest cost over value and fail to capture longer-term or system-wide benefits.
Workforce capacity and readiness	Adoption means training the existing health-care workforce to use, interpret and integrate tools into clinical workflows, often within overstretched systems. Resistance may arise when technology is perceived as a threat to job security or professional autonomy.
Social acceptance and trust	Public perception of AI errors differs from human errors, with AI mistakes in some cases weighed more heavily, even when performance is comparable or superior. Trust depends on transparency, accountability, and overcoming social-acceptance challenges, as patients can be reluctant to view AI as a useful tool, especially in patient-facing roles.
Regulatory misalignment	Regulatory pathways might be slow, unclear and/or designed for traditional devices rather than adaptive AI systems. Fragmented regimes across jurisdictions can further delay deployment and increase costs.
Bias and inequality	Many AI systems are trained on limited high-income country data sets with different disease profiles, demographics and

documentation practices, limiting transferability across contexts. Performance may degrade with differences in epidemiology, data quality and population characteristics, risking inequitable outcomes. Limited language coverage and localisation can further constrain use and exacerbate inequalities. The benefits of AI may not reach populations without the infrastructure and skills needed to use it being in place.

Source: TBI

These barriers are among the most notable, but they are not exhaustive. Others include liability and legal uncertainty around accountability when AI systems fail and vendor dependency and sustainability concerns when external support ends, as well as misaligned incentives (including for incumbent providers or interest groups who benefit from maintaining current arrangements). The point is not to catalogue every obstacle, but to recognise that evidence of clinical efficacy, while necessary, is not sufficient for adoption.

A technology that is clinically sound, cost-effective and superior to current practice may still fail to diffuse if the conditions for implementation are not in place; as with any clinical intervention, it will still require robust governance to minimise any potential adverse effects. Addressing these barriers requires coordination across ministries, alignment of incentives and investment in the enabling infrastructure (technical, financial, institutional and social) that makes adoption of the right tools feasible.²⁷ This is particularly urgent in settings where the gap between need and capacity is widest, and where the potential returns from successful adoption are correspondingly greatest.

05

Recommendations

Despite the barriers mentioned above, there are actions that governments can take now to proactively build the conditions for successful and ethical adoption. The challenge is not just to identify the right tools, but to implement them in ways that improve services and outcomes. Governments can take action now to start realising the benefits of AI in the short term: they can implement solutions that work within current systems while building the conditions for further adoption in the future.

1. Reframe Regulation Around Comparative Risk

Regulatory approaches should differentiate clearly between technologies with different levels of risk and potential impact, and explicitly take account of the risk of sticking with the status quo. This requires new mechanisms that are designed for emerging technologies, with approval processes proportionate to risk, better quantification of pre-existing risk and a shift beyond a predominantly pre-market model towards greater emphasis on post-market surveillance, including assessing for bias and impact on inequalities. Testing and monitoring AI tools in the contexts where they will be used will make them more locally relevant, representative and equitable.

Regulation and decisions around efficacy of technologies should remain within regulatory bodies that operate at arm's length from government. This would maintain objective and consistent assessment of risks and benefits, and help to depoliticise adoption.

2. Set a Clear Direction and Communicate Openly

Governments that introduce AI into health systems without explaining why will quickly find that public trust, once lost, is hard to regain. They must therefore be explicit not only about the benefits of adoption but also the harms that are occurring in the current system: delayed diagnoses, missed conditions and preventable deterioration. Governments should articulate where and why technology will be used, the benefits it is expected to

deliver, how evidence-based decisions will be made on what is introduced and how risks will be minimised after introduction. Done well, this also helps to manage political risk by making trade-offs explicit and building public understanding of rationale.

3. Act Decisively on High-Impact Priorities

A sequenced and costed implementation plan is required, with clear milestones that move beyond pilots and set out a long-term approach to improving health care, underpinned by adequate financing and clear accountability for delivery, including infrastructure development. This means clearly prioritising tools that address real health-system challenges (such as capacity constraints, diagnostic gaps and failures in care coordination), even where this involves accepting and carefully governing different levels of risk in pursuit of greater impact, rather than defaulting to lower-risk, lower-impact solutions. It is important to proceed with implementation in ways that strengthen (rather than fragment) the system, and learn lessons for future implementation.

4. Enable Clinicians, and Co-Design With Users

Alongside regulatory clarity, governments should actively enable adoption by prioritising tools that work for clinicians and patients in practice. Evidence from human-centred policy approaches shows that systems are more likely to succeed when they are designed with users and not just for them, shaped through early engagement, co-design and prototyping, as well as iteration based on real-world feedback. This means working with health-care workers to prioritise technologies that integrate with existing workflows and demonstrably reduce burden, as these are far more likely to be used in practice. It also requires dedicated funding and protected time for workforce training and public engagement, including ongoing, role-specific support to build confidence and capability over time.

5. Build the Foundations and Environment for Adoption

Effective deployment depends on core capabilities, including data security, interoperability and digital connectivity. Some applications will struggle to function in their absence, so governments need to strengthen foundations and increase the range of tools that can be deployed. They should also make use of tools that can work within their existing infrastructure, as this will deliver early benefits, build momentum and enable learning for subsequent deployments.

A clear and credible plan can attract innovation. In its absence, and where regulatory and procurement pathways are opaque, technology companies are often unable to enter the market at all, meaning relevant solutions are neither developed nor deployed locally. Governments need to be clear on what tools they need, what funding is available and the steps that private-sector organisations need to go through for their tools to be deployed.

06

Conclusion

Evidence of the impact of AI in health is growing by the day; it has the potential to bridge huge capacity gaps in health systems around the world. But governments and decision-makers are hesitant, in part because of natural caution about introducing new technologies in a field that deals with life and death. In practice, this often leads to the risks of AI adoption being given greater weight than the risks of delay – treating inaction as the safer option, even where it is not.

Decisions should be based on a comparison of impact and incremental risk that is relative to current practice, with deployment focused on the most pressing health-system constraints rather than the technologies themselves. Overcoming these constraints means taking managed risks, backed by appropriate safeguards. In doing so, governments have the opportunity to expand access, improve quality and save lives.

07

Acknowledgements

The authors would like to thank the following experts for their input and feedback (while noting that contribution does not equal endorsement of points made in the paper).

Frank Anthony, Ministry of Health Guyana

Dame Sally Davies, Trinity College, Cambridge

Katie Drasser, RockHealth.org

Tom Fairfax, University of Sunderland

Navid Farzad, Frist Cressey Ventures

Johann Go, University College London

Anurag Mairal, Stanford Biodesign

Greg McDavitt, ServiceNow

Charlotte Rock, General Catalyst Institute

Lawrence Tallon, Medicines and Healthcare products Regulatory Agency

Endnotes

- 1 <https://www.thelancet.com/journals/lancet/article/PIIS0140-6736%2818%2931668-4/fulltext>
- 2 [https://www.thelancet.com/journals/lancet/article/PIIS0140-6736\(24\)01439-9/abstract](https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(24)01439-9/abstract)
- 3 <https://www.england.nhs.uk/long-read/monthly-operational-statistics-june-2025/>
- 4 <https://www.england.nhs.uk/2026/02/waiting-list-lowest-3-years-battled-busiest-winter-on-record/>
- 5 <https://www.england.nhs.uk/2026/05/health-service-hits-18-week-target-amid-half-million-waiting-list-drop/>
- 6 <https://www.ft.com/content/cb763b6e-3c1f-4c9a-9fdb-5188fb102f68>
- 7 <https://kmpdc.go.ke/Registers/medical%5Fspecialist%5Fpractitioners.php>
- 8 <https://data.worldbank.org/indicator/SP.POP.TOTL>
- 9 Estimation based on the assumed number of radiologists required in Kenya, calculated using the average number of radiologists per population across OECD countries as a benchmark:
<https://pmc.ncbi.nlm.nih.gov/articles/PMC8613895/>
- 10 <https://www.who.int/teams/health-workforce/hwf%5Fgovernance/health-workforce>
- 11 <https://www.who.int/data/gho/data/indicators/indicator-details/GHO/community-health-workers-%28number%29>
- 12 <https://www.who.int/docs/default-source/documents/almaata-declaration-en.pdf>
- 13 <https://pmc.ncbi.nlm.nih.gov/articles/PMC6701148/>
- 14 <https://www.thelancet.com/journals/lancet/article/PIIS0140-6736%2808%2961405-1/fulltext>
- 15 <https://www.nature.com/articles/s41591-018-0107-6.epdf?>
- 16 <https://www.yana.ai/>
- 17 <https://www.who.int/news-room/fact-sheets/detail/sepsis>
- 18 <https://www.nature.com/articles/s41746-023-00986-6>

- 19 <https://www.nature.com/articles/s41467-021-20910-4>
- 20 <https://www.yana.ai/>
- 21 <https://www.sciencedirect.com/science/article/pii/S0959804919302217>
- 22 <https://www.thelancet.com/journals/lancet/article/PIIS0140-6736%2825%2902464-X/abstract?>
- 23 <https://jamanetwork.com/journals/jama/article-abstract/2760710>
- 24 <https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/2771093> in comparison to <https://www.sciencedirect.com/science/article/abs/pii/S0140673619317210>
- 25 See, among others, Cass R. Sunstein, *Laws of Fear: Beyond the Precautionary Principle* (Cambridge University Press, 2005); John D. Graham and Jonathan B. Wiener (eds.), *Risk vs. Risk: Tradeoffs in Protecting Health and the Environment* (Harvard University Press, 1995); Jonathan B. Wiener and John D. Graham, “The Precautionary Principle and Risk-Risk Tradeoffs: A Comment”, *Journal of Risk Research* 11, no. 4 (2008): 465–74
- 26 <https://www.nature.com/articles/s41746-024-01213-6>
- 27 <https://dalberg.com/wp-content/uploads/2016/12/ReadySetLaunch%5FDec23%5FWeb-508.pdf>

Follow us

facebook.com/instituteglobal

x.com/instituteGC

instagram.com/institutegc

General enquiries

info@institute.global

Copyright © May 2026 by the Tony Blair Institute for Global Change

All rights reserved. Citation, reproduction and or translation of this publication, in whole or in part, for educational or other non-commercial purposes is authorised provided the source is fully acknowledged Tony Blair Institute, trading as Tony Blair Institute for Global Change, is a company limited by guarantee registered in England and Wales (registered company number: 10505963) whose registered office is One Bartholomew Close, London, EC1A 7BL.