The Post-Pandemic Risk in Health AI

A two-page memo for leaders responsible for deploying clinical AI

By Dr Philip McMillan – Physician | Researcher

The core assumption that no longer holds

Most Health AI systems are being deployed on the assumption that population biology remains broadly stable. That assumption underpins model training, label validity, performance evaluation, and safety claims. Since 2020, it has become increasingly unreliable.

What changed at population level after the pandemic

The pandemic introduced a population-scale biological shift characterised by immune dysregulation, endothelial and microvascular injury, autonomic instability, and prolonged post-viral inflammatory states. These changes are heterogeneous, often subclinical, and unevenly distributed across age groups and comorbidities. They do not present as new diseases, but as altered baselines.

Why this matters clinically, not politically

In clinical practice, post-pandemic biology manifests as broader phenotypes, overlapping syndromes, slower recovery trajectories, and increased diagnostic ambiguity. Presentations are more complex, less specific, and more trajectory dependent. Viewed through Health AI data, this is considered noise, but from a biological standpoint, it indicates a significant shift.

Why most AI systems misinterpret the signal

Health AI systems are mainly designed to reduce variance, ensure diagnostic accuracy, and perform well against historical data. Post-pandemic biology violates these design assumptions. Labels break down, syndromes overlap, and long-term progress is more important than isolated moments. Supervised models are prone to label instability, making confidence calibration unreliable.

Why validation no longer protects you

Traditional validation checks generalisation within set distributions, but it fails to assess safety during biological changes. Metrics can appear reliable even as clinical performance declines. Most failures happen as unnoticed misclassifications or unwarranted confidence.

The Post-Pandemic Risk in Health AI

A two-page memo for leaders responsible for deploying clinical AI

By Dr Philip McMillan – Physician | Researcher

Why I recognised this early

My perspective comes from working clinically before, during, and after the pandemic, and from analysing post-viral and post-COVID patterns as evolving biological phenomena rather than isolated conditions. At the pandemic's start, I prioritized immune and vascular factors over respiratory labels, letting me identify key trends before official recognition. A similar approach is now being utilised to assess risks associated with Health AI.

The missing layer in current Health AI

Many AI stacks do not include an inference layer that can identify phenotype drift, manage label uncertainty, or reason across syndromes and trajectories. In the absence of such a layer, biological instability following the pandemic is often interpreted as error rather than valuable signal, and uncertainty is minimized instead of being highlighted. This represents a conceptual gap rather than a computational limitation.

Strategic implications for Health AI leaders

Heavy reliance on pre-2020 data is now a liability. Label quality should be questioned, not taken for granted. Focusing on syndromes or trajectories is often safer than making definitive diagnoses. Systems that do not explicitly account for post-pandemic biology will underperform in complex, high-risk patients and expose organisations to clinical, ethical, and reputational risk.

What needs to happen next

Health AI should audit biological assumptions, test models for post-pandemic phenotypes, and include clinicians focused on mechanisms and disease progression. This is not a call to slow AI adoption, but to prevent avoidable failure as systems scale.

Call to action

If you are building, deploying, regulating, or investing in Health AI and want an independent assessment of post-pandemic biological risk, request a briefing or strategic review.