Initial provider resistance to combination treatments was probably less of a barrier for what was then the relatively new problem of HIV infection, compared to CVD for which decades-old treatment paradigms are deeply ingrained. Use of combination therapy for HIV/AIDs was also driven by recognition of the therapeutic emergency posed by HIV, which facilitated paradigm shifts, regulatory approvals and essential medicines listings, all of which led to the rapid scale up of affordable fixed-dose combination treatment for individuals with HIV infection. These drivers of uptake do not exist for CVD.

Since the concept was first discussed just over 20 years ago, there has been substantial activity relating to CVD polypill development, research and advocacy (Fig. 1). But scale up remains elusive. Encouragingly, some progress has been made with polypills for hypertension, with the WHO Model List of Essential Medicines listing dual-combination blood-pressure-lowering drugs for initial treatment in 2019, followed more recently by a matching recommendation in updated WHO hypertension guidelines[23,24]. Although this may herald a trend toward promoting polypill-based approaches, realizing their potential will only happen with a global shift in treatment paradigms, new business models and solutions to implementation challenges. This in turn requires urgent consensus building among consumers, providers, payers, manufacturers and a range of other major private and public stakeholders. The risks of delay might be another 20 years before any meaningful progress occurs, at the cost of countless avoidable premature deaths globally. ❐

Anushka Patel[1]✉, Dike Ojji[2], H. Asita de Silva[3], Stephen MacMahon[1,4] and Anthony Rodgers[1,4]

[1]*The George Institute for Global Health, University of NSW, Sydney, NSW, Australia.* [2]*Department of Medicine, University of Abuja, Abuja, Nigeria.* [3]*Clinical Trials Unit, Department of Pharmacology, Faculty of Medicine, University of Kelaniya, Kelaniya, Sri Lanka.* [4]*School of Public Health, Imperial College London, London, UK.*

✉*e-mail: apatel@georgeinstitute.org*

Published online: 31 January 2022
https://doi.org/10.1038/s41591-021-01635-9

References

1. Roth, G. A. et al. *J. Am. Coll. Cardiol.* **76**, 2982–3021 (2020).
2. Yusuf, S. et al. *Lancet* **378**, 1231–1243 (2011).
3. Peiris, D. et al. *PLoS Med.* **18**, e1003485 (2021).
4. NCD Risk Factor Collaboration. *Lancet* **398**, 957–980 (2021).
5. Webster, R. et al. *Int. J. Cardiol.* **15**, 147–156 (2016).
6. Castellano, J. M. et al. *J. Am. Coll. Cardiol.* **64**, 2071–2082 (2014).
7. Muñoz, D. et al. *N. Engl. J. Med.* **381**, 1114–1123 (2019).
8. Webster, R. et al. *JAMA* **320**, 566–579 (2018).
9. Chow, C. K. et al. *Lancet* **398**, 1043–1052 (2021).
10. Lin, J. K. et al. *Lancet Glob. Health* **7**, e1346–e1358 (2019).
11. Lung, T. L. et al. *Lancet Glob. Health* **7**, e1359–e1366 (2019).
12. Gaziano, T. A. et al. *Am. Heart J.* **214**, 77–87 (2019).
13. Liu, H. et al. *Circ. Cardiovasc. Qual. Outcomes* **8**, 301–308 (2015).
14. Wood, F. et al. *BMJ Open* **5**, e008018 (2015).
15. Cosin-Sales, J. et al. *J. Comp. Eff. Res.* **10**, 975–985 (2021).
16. Wang, N. et al. *Int. J. Cardiol.* **341**, 96–99 (2021).
17. Roshandel, G. et al. *Lancet* **394**, 672–683 (2019).
18. Yusuf, S. et al. *N. Engl. J. Med.* **384**, 216–228 (2021).
19. Joseph, P. et al. *Lancet* **398**, 1133–1146 (2021).
20. Kishore, S. P. et al. *J. Am. Coll. Cardiol.* **71**, 564–574 (2018).
21. Khatib, R. et al. *Lancet* **387**, 61–69 (2016).
22. Do, N. T. et al. *BMJ Glob. Health* **6**, e006523 (2021).
23. Salam, A. et al. *J. Clin. Hypertens.* **22**, 1769–1779 (2020).
24. WHO. *Guideline for the Pharmacological Treatment of Hypertension in Adults* (2021).
25. Wald, N. J. & Law, M. R. *BMJ* **326**, 1419 (2003).

Author contributions

A.P. wrote the first draft. All other authors made critical revisions.

Competing interests

A.P., S.M. and A.R. are employed by The George Institute for Global Health (TGI), which holds patents for ultra-low-dose-fixed combination products for the treatment of hypertension and diabetes. S.M. and A.R. are listed as inventors on these patents. S.M. is a Director on the boards of George Health Enterprises Pty Ltd (GHE) and the social enterprise arm of TGI, as well as its subsidiary George Medicines Pty Ltd (GM); these companies have received funding from public and private investors to conduct trials of fixed combination products for regulatory approval. A.R. is seconded part-time to GM. A.P., S.M. and A.R. do not have direct financial interests in any of the patent applications or investments. D.O. and H.A.d.S. declare no competing interests.

Clinical decisions using AI must consider patient values

Built-in decision thresholds for AI diagnostics are ethically problematic, as patients may differ in their attitudes about the risk of false-positive and false-negative results, which will require that clinicians assess patient values.

Jonathan Birch, Kathleen A. Creel, Abhinav K. Jha and Anya Plutynski

Recent years have seen a surge of interest in medical applications of machine learning and artificial intelligence (AI) — and in the ethics of these applications[1,2]. In the near future, it is unlikely that AI technology will replace human decision-makers, but it is likely to assist human decision-makers in many contexts[3]. For example, machine learning can assist with the diagnosis of cancer from imaging data[4], the mapping of the boundaries of tumors[5], the treatment of sepsis[6], the classification of field-based rapid HIV tests[7] and the diagnosis of cognitive-motor dissociation in unresponsive patients with brain injuries[8].

Probabilistic classifiers

Many AI methods with strong promise in medical imaging are probabilistic classifiers: their native output is a probability. This includes naive Bayes classifiers, decision trees and multiple neural-network-based approaches[9]. For example, a probabilistic classifier for diagnosing cancer from oncological image data will generate a probability that cancer is present. This leads to an important design 'choice point'. The software provided to clinicians may hide the raw probabilities and build in decision thresholds for converting probabilities into recommendations, so that a probability of cancer above 70% (for instance) gives the output 'recall recommended'. Alternatively, the software may provide the raw probabilities, with the burden falling on clinicians to take the probabilities into account when making decisions. A third possibility is that the software might incorporate information about the patient's own attitudes about risk and the value (positive or negative) that they assign to different possible outcomes, together with the probabilities estimated by the algorithm, so as to generate a recommendation tailored to this particular patient.

The first option may seem appealing on the grounds of simplicity and ease of use. However, there are

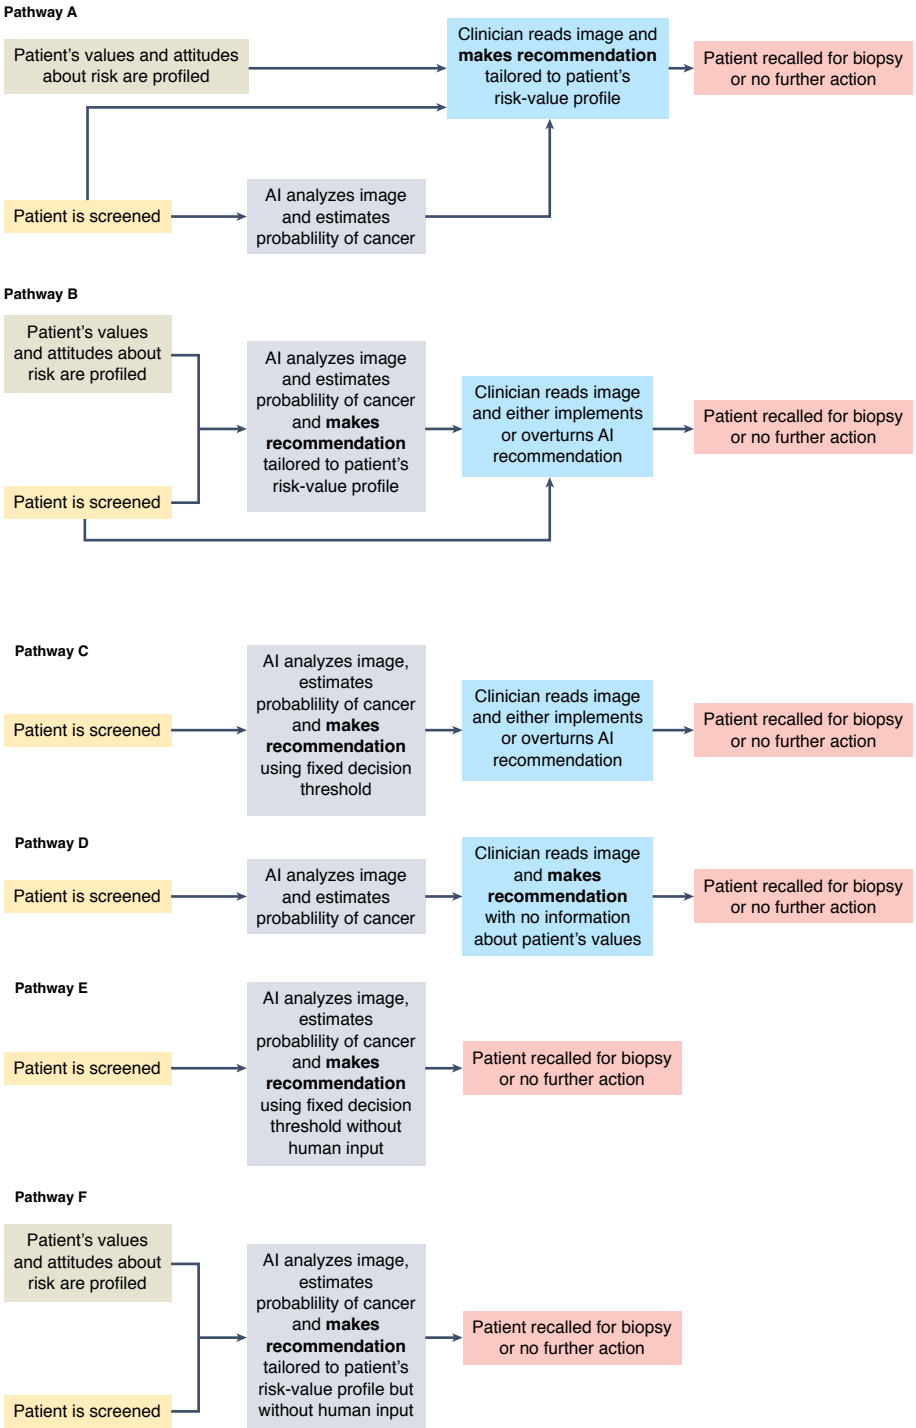

Fig. 1 | Six possible decision pathways using AI algorithms for breast cancer screening. A summary of the ethical implications of the different pathways is presented in Table 1. We recommend further research aimed at realizing pathways A and B, but have notable concerns about pathways C, D, E and F.

ethical reasons that developers of medical AI should take one of the other two options. Clinical decision-makers should be provided either with probabilistic outputs or with a recommendation that takes both the probabilities and the patient's values and tolerance for risk into account. This is because in clinical settings, there can be no one-size-fits-all decision threshold. From an ethical point of view, it is appropriate for the decision threshold to be sensitive to the values and attitudes toward risk that this particular patient holds.

Decision thresholds

The case against built-in decision thresholds can be made with the example of breast cancer screening. When a human reads a mammogram as part of an initial screening protocol, their output is a decision: recall the patient (or not) for

Table 1 | An ethical analysis of possible decision pathways

Decision pathway from Fig. 1	Sensitive to patients' values and preferences?	Likely to have patient support?	Concerns
A	Yes	Yes	Implementation challenges (discussed in main text)
B	Yes	Yes	Implementation challenges (discussed in main text)
C	No	No	Despite human input, patients' values and preferences are not considered when managing uncertainty, eroding patient trust
D	No	No	Despite human input, patients' values and preferences are not considered when managing uncertainty, eroding patient trust
E	No	No	Loss of human input, combined with failure to consider patients' values and preferences, is likely to severely undermine trust
F	Yes	No	Although patients' values and preferences are considered, loss of human input is likely to undermine patient trust at the present time

further investigation, typically a biopsy. By contrast, the output of several AI systems is a "continuous score that represents the likelihood of cancer being present"[4]. This probability from the AI algorithm can be converted to a decision in a range of ways, depending on where (between 0% and 100%) the critical threshold is set, which triggers a recall decision.

Different decision thresholds can be imposed for different purposes. A lower threshold could lead to more false-positive diagnoses, whereas a higher threshold could lead to more false-negative diagnoses, whereby cancer is missed. Choosing the decision threshold requires weighing of the disvalue for each kind of error, which in turn requires context-sensitive value judgements[10–12].

The programmer of the algorithm could choose a single threshold, set at the same level for all patients. This would be suboptimal for patient autonomy, as the patient's own values would be ignored[13]. Instead, the algorithm could produce a raw continuous score, leaving the threshold judgement between the clinician and the patient (pathway A in Fig. 1). Alternatively, the algorithm could take account of the patient's own values and risk tolerances when recommending a threshold, perhaps by calculating a utility measure that takes risk tolerance into account[14] (pathway B in Fig. 1).

Risk-value profiles

There are various ways to acquire information about a patient's values. One would be for patients, at the time of screening or earlier, to be given a risk-profiling questionnaire that probes their values and attitudes about different outcomes, which would allow the construction of a risk-value profile for the patient. This idea is inspired by the risk-profiling questionnaires commonly used in finance. These questionnaires typically ask investors for their level of agreement with statements about risk, such as "Generally, I prefer investments with little or no fluctuation in value, and I'm willing to accept the lower return associated with these investments"[15].

A risk-profiling questionnaire suitable for cancer screening would probe the patient's attitudes about the risk of over-diagnosis, false-positive and false-negative results, and over-treatment versus under-treatment, and the expected value to the patient of additional years of life of varying quality levels. The questionnaire might also ask patients to respond to statements such as 'I would rather risk surgical complications to treat a benign tumor than risk missing a cancerous tumor'.

Patient support

Recently, a large study investigating the Dutch population's view on the use of AI for the diagnostic interpretation of screening mammograms[16] found that the general population currently does not support fully independent use of AI without the involvement of a human radiologist as well. This suggests that it would not be appropriate, at the present time, to pursue a decision pathway in which there is no role for a human reader, even if the AI does take the patient's values into account (as in pathways E and F in Fig. 1).

In principle, the attitudes of patients, who have relevant lived experience and a personal stake in the matter, could differ from those of the general population. However, a survey of United States–based patient-advocacy groups, carried out by the authors of this Comment (with institutional review board approval from Washington University in St Louis), also found resistance to clinical decisions being made without human input[17]. Most respondents were comfortable with some involvement for AI, but only when there was also a human reader. Feedback included such comments as "I see AI as a tool to assist clinicians in medical decisions. I do not see it as being able to make decisions that effectively weigh my personal input or really have the clinical experience and intuition of a good physician." Another respondent commented that AI "could be a valuable tool, but combined with physicians' expertise and consultation with patients." There were many other comments along the same lines.

In addition to their concerns about the potential loss of human input with the advent of AI, most respondents were concerned by the thought that uncertainty about their screening image might be hidden from them, either by a human reader or by an algorithm applying a decision threshold. Moreover, most were concerned by the prospect of the uncertainty being managed (either by a human reader or by an algorithm) without any consideration of their values and preferences (as in pathways C and D in Fig. 1). One of the risks of increased use of AI is that these fears (of uncertainty being hidden, and of one's values being ignored) will become more acute. But there is also a corresponding opportunity, because AI allows the possibility of varying decision thresholds in a transparent way that demonstrates sensitivity to these concerns.

Implementation challenges

There are a number of implementation challenges with the proposed decision pathways A and B (Fig. 1) that will require further research. First, the design of any survey to elicit a risk-value profile would require consultation and trialing, given the potential for framing effects to influence the patient's answers. A good design would also build in appropriate over-rides. Returning to the financial analogy, a good risk-profiling questionnaire will allow an investor to communicate that even though they usually have an appetite for high risk, they want to be cautious on this particular occasion (or vice versa). Similarly, a good questionnaire for probing patients' values would allow a patient to

state that even though they would usually assign great disvalue to a false-positive result, their priority on this occasion is to avoid a false-negative result (or vice versa). Moreover, patients must be properly informed about the relevant concepts. Many patients are unfamiliar with the concept of over-diagnosis and therefore may be unable to weigh the relative risk of unnecessary diagnosis and treatment against the risk of failing to discover a cancer[18]. Moreover, patients may not always have preferences about such outcomes. There must still be a sensible default decision threshold that can be used in cases in which patients choose to withhold their attitudes or simply have no preferences.

There is also a danger of exaggerating the precision of the probabilities. If the dataset used to train the algorithm was small or non-representative, a probability range may be a more reasonable output than a precise probability. There is also a risk that clinicians will be unwilling or unprepared to take the patient's risk-value profile into account. These recommendations would create a more complex decision task for clinicians than reliance on a pre-programmed threshold; this emphasizes the importance of training in the use of AI in clinical settings and the co-design of diagnostic devices with physicians and patients.

These challenges notwithstanding, tailoring decision thresholds to the patient through the use of information about the patient's values and attitude about risk is vastly preferable to leaving them to be fixed by the software developer in a one-size-fits-all manner. ❐

Jonathan Birch[1]✉, Kathleen A. Creel[2], Abhinav K. Jha[3] and Anya Plutynski[4]

[1]*Centre for Philosophy of Natural and Social Science, London School of Economics and Political Science, London, UK.* [2]*McCoy Family Center for Ethics in Society and the Institute for Human-Centered Artificial Intelligence, Stanford University, Stanford, CA, USA.* [3]*Department of Biomedical Engineering and Mallinckrodt Institute of Radiology, Washington University in St. Louis, St. Louis, MO, USA.* [4]*Department of Philosophy and Division of Biology and Biomedical Sciences, Washington University in St Louis, St Louis, MO, USA.*
✉*e-mail: j.birch2@lse.ac.uk*

Published online: 31 January 2022
https://doi.org/10.1038/s41591-021-01624-y

References

1. McCradden, M. D. et al. *Nat. Med.* **26**, 1325–1326 (2020).
2. Grote, T. & Berens, P. *J. Med. Ethics* **46**, 205–211 (2019).
3. Tschandl, P. et al. *Nat. Med.* **26**, 1229–1234 (2020).
4. McKinney, S. M. et al. *Nature* **577**, 89–94 (2020).
5. Liu, Z. et al. *Phys. Med. Biol.* **66**, 124002 (2021).
6. Komorowski, M. et al. *Nat. Med.* **24**, 1716–1720 (2018).
7. Turbé, V. et al. *Nat. Med.* **27**, 1165–1170 (2021).
8. Claassen, J. et al. *N. Engl. J. Med.* **380**, 2497–2505 (2019).
9. Hastie, T. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Springer, 2009).
10. Plutynski, A. in *Exploring Inductive Risk: Case Studies of Values in Science* (eds. Elliott, K. C. & Richards, T.) 149–170 (Oxford University Press, 2017).
11. Douglas, H.E. *Science, Policy, and the Value-free Ideal* (University of Pittsburgh Press, 2009).
12. Bright, L. K. *Synthese* **195**, 2227–2245 (2017).
13. Wenner, D. M. *Int. J. Fem. Approaches Bioeth.* **13**, 28–48 (2020).
14. Buchak, L. *J. Med. Ethics* **43**, 90–95 (2016).
15. Pan, C. H. & Statman, M. *J. Invest. Consult.* **13**, 54–63 (2012).
16. Ongena, Y. P. et al. *J. Am. Coll. Radiol.* **18**, 79–86 (2021).
17. Birch, J., Creel, K., Jha, A. & Plutynski, A. *Zenodo* https://doi.org/10.5281/zenodo.5589207 (2021).
18. Nagler, R. H. et al. *Med. Care* **55**, 879–885 (2017).

Acknowledgements

We thank S. Bhalla, L. Kofi Bright, A. Houston, L. Hudetz, R. Short, J. Swamidass, K. Vredenburgh, Z. Ward, K. Wright and patient groups at Washington University in St Louis, Stanford University and Johns Hopkins University for their input and advice. A.K.J. acknowledges support from the National Institute of Biomedical Imaging and Bioengineering of the US National Institutes of Health (R01-EB031051 and R56-EB028287).

Author contributions

All authors contributed to conceptualization, methodology (survey design), investigation (consulting patients), and writing (review and editing). J.B. wrote the original draft.

Competing interests

The authors declare no competing interests.

Check for updates

The French health pass holds lessons for mandatory COVID-19 vaccination

The passe sanitaire increased levels of vaccination, but to a lower extent among the most vulnerable, and did not reduce vaccine hesitancy itself, showing the importance of outreach to underserved communities and the potential limits of mandatory vaccination policies.

Jeremy K. Ward, Fatima Gauna, Amandine Gagneux-Brunon, Elisabeth Botelho-Nevers, Jean-Luc Cracowski, Charles Khouri, Odile Launay, Pierre Verger and Patrick Peretti-Watel

Public authorities in many countries are considering mandating vaccination against COVID-19 for the whole eligible population[1]. Most countries are confronted with the difficulties of reaching the vaccination rates obtained for diseases such as measles, which are often above 95%. During the summer of 2021, French authorities implemented a health pass, or passe sanitaire, requiring everyone aged 12 and older to present proof of vaccination or a negative test for SARS-CoV-2 to access a wide array of public spaces, including bars, libraries and hospitals. The introduction of the passe sanitaire markedly increased the number of people vaccinated against COVID-19. But, as of November 2021, coverage is plateauing at around 90% of the eligible population and a debate has arisen on whether the next step should be mandating this vaccination[2].

There are lessons to be learnt from the French experience with the health pass that contribute to the current debate on mandatory COVID-19 vaccination.

Barriers to vaccination

In France, vaccination coverage against COVID-19 rose steadily during the first half of 2021 until it reached a first plateau in mid-June, with around 60% of the adult population having had at least one dose[2]. After having vaccinated the most willing, public authorities were confronted with three classic barriers to vaccination[3] (Table 1). The first was doubts regarding the safety of vaccines against COVID-19[4,5]. The proportion of the population who intended to receive the