Framing Tomorrow: Who Decides What is Ethical in Foresight?

Dubai Future Forum 2024 Learnings Day

Facilitation reflections report

Introduction

The Nuffield Council on Bioethics (NCOB) is working to put ethics at the centre of decisions regarding biomedicine and health, so we all benefit. We believe by bringing ethical considerations more overtly into policy development we can create governance that better reflects public values and priorities, and by doing this, policies can garner a higher degree of public trust.

Exploring and examining the ethical impacts of emerging technologies in the biomedicine and health fields has received renewed interest in recent years, particularly given its crucial role in ensuring responsible research and innovation. Since the COVID-19 pandemic, preparedness, strategic foresight and intelligence, anticipatory governance and long (or long*er*)-term thinking have been increasingly prioritised in health and care policy. However, scoping conducted by the NCOB identified a gap in how current foresight approaches in the UK healthcare innovation policy context identify and address emerging ethical considerations. For example, the potential harms to vulnerable groups such as discrimination, exclusion or stigmatisation. Developing and incorporating horizon scanning, futures and foresight methods that can enable better mitigation of these issues would help ensure decisions on how technological innovations are utilised is going to benefit all and not simply the lucky few.

In an ambition to address this gap and bring these issues to the policy table, the NCOB has embarked on an ambitious project to develop a range of foresight tools and approaches that will assist policymakers in their efforts to explore ethical implications at an earlier stage of the innovation process. As part of the NCOB's ethical lens for horizon scanning and foresight project, we are keen to ensure that ethical foresight is routinely considered as part of and certainly, alongside existing, common-place methods.

Purpose of this report

This report summarises the outcomes of an experimental workshop designed and facilitated by the NCOB, the Ada Lovelace Institute and Associate Professor Federica Lucivero (Ethox Centre, University of Oxford) on the 2024 Dubai Future Forum Learnings Day.

The experiment invited 20 conference attendees to explore the impact of ethical foresight in decision-making. Participants were divided into two groups: one engaged with traditional scenario-based discussions, and the other worked with ethically enriched scenarios featuring moral tensions and explicit trade-offs. This division aimed to assess how different scenario approaches influenced the articulation of values, ethical awareness, and resultant policy recommendations.

Workshop objectives:

1. Pilot ethically enriched techno-moral scenarios as a means to enhance ethical input in/into decision-making.

2. Compare this approach with traditional methods to identify their respective impacts on participant engagement, value articulation, and handling of ethical differences.

Workshop structure

Introduction (10mins)

Participants were given information about the facilitators, explaining the organisations and their purpose. They were then told they were taking part in an experiment and that more would be revealed later on.

Parallel group sessions (120mins)

Group 1: Traditional scenario discussions focusing on neutral, operational contexts of genomic health prediction (GHP) use.

Group 2: Ethically enriched scenarios with narrative elements highlighting moral dilemmas and stakeholder conflicts.

Reveal and discussion (40mins)

The groups were reconvened and told what the experiment had been testing and why. They were then asked to share their feedback on how things felt as participants and to engage in a comparative discussion of the ethical issues raised.

Workshop scenarios

Group one: Traditional scenarios (Annex 1)

Participants discussed two scenarios that were grounded in plausible futures for GHP.

- Widespread prevention model (2030): GHP used sparingly, under strict regulation, to improve treatment outcomes for high-risk cases. In this scenario, public trust was high due to robust data protection and governance standards.

- Broad public-use model (2030): Deregulated GHP targeting the majority of the population, integrating AI systems and private insurers and public concerns about data misuse and insurance discrimination were prominent.

Group two: Techno-ethical scenarios (Annex 2)

Participants were presented with complex moral tension through the aid of techno-moral scenarios.

- Patient-focused versus doctor-mediated models: A patient-focused approach emphasised individual autonomy in health and care management and a doctor-guided model highlighted professional expertise in mediating health and care decisions.

- Automation versus human judgement: A growing reliance on GHP systems and the potential erosion of clinician judgment – balance between efficiency, safety, and the risks of automation bias.

Workshop reflections

Both groups emphasised the critical importance of transparent governance to build and maintain public trust in GHP. They also frequently debated the tension between cost-saving

measures and equitable access to healthcare. And explored concerns about over-reliance on GHP technologies and the potential erosion of clinician judgment.

However, there were some differences seen across the groups with the ethically enriched scenarios appearing to facilitate a richer, more nuanced discussion.

Group two participants were more engaged in exploring moral tensions and articulating their positions, frequently debating trade-offs between competing values like autonomy and paternalism. Group two also demonstrated greater awareness of the long-term societal impacts of GHP, particularly regarding data privacy, equity, and the risks of over-automation.

Group one's discussions tended to focus on feasibility and technical details. It highlighted trust and regulatory issues but often omitted broader ethical considerations, focusing on practical resolutions without delving deeply into value conflicts.

The following table highlights key differences between the groups:

Aspect	**Group one (control)**	**Group two (ethical vignettes)**
Articulation of values	This was somewhat limited; the predominant focus was on operational issues	Explicit: linked to societal, political, cultural and legal values
Awareness of ethics and autonomy	Focused on trust and governance	Broader focus, including equity
Handling disagreements	Consensus-driven	Explored and examined the posed moral disagreements
Long-term impact awareness	Fewer discussions about the long-term consequences and impact of the scenario.	More discussions about what the future could look like in different jurisdictions.

Workshop participant feedback

Group one

Participants in group one, our control group, reported that they felt the scenarios were more factual and less dramatic. This meant that when implications were discussed, they felt much more theoretical and this could have impacted on the level of empathy felt.

Ethical issues were explored, but some suggested this was likely because the group knew the premise of the workshop was around ethical foresight based upon the conference title.

The group suggested that the lack of examples in the scenarios led to them introducing some and this could have biased the conversations. Moreover, when asked about what was 'good' and what was 'bad', participants felt this was too vague.

The group were satisfied with the number of participants in their group but felt that the number of questions from the facilitators meant that there was a potential dilution from the focus.

Group two

Participants in group two, our experiment group, reported that the presentation of the scenarios was able to resonate with them across their different national contexts, and that this allowed them to discuss issues of choice and agency more deeply.

Participants felt that the scenarios effectively highlighted ethical quandaries. They also illustrated the complexity of current decision-making processes, demonstrating the importance of not relying on a single view.

The detailed nature of the scenarios got mixed responses from the workshop participants. Some felt they helped them to understand the operational environment and that this facilitated discussions about uncertainty. However, others argued the detailed nature was somewhat leading, particularly when it came to policy solutions that seemed predetermined. And it was noted that too much detail could hinder rather than enhance the conversation. This led to questions about how to best "signpost but not lead," emphasising the importance of open exploration.

All reported that including drama in the scenarios worked well for fostering future-focused conversations and helped draw out discussions about values. However, the diversity of perspectives meant that in-depth ethical discussions were limited due to the time constraints of the workshop.

Our conclusion

We believe the workshop demonstrated a value for ethically enriching scenarios in a future focused workshop.

Our observations and participant feedback suggest that it proved to be effective in eliciting deeper engagement with issues of moral responsibility, value pluralism and fairness, revealing a broader range of ethical issues compared to traditional approaches. These findings support a case for adopting such methodologies as they may strengthen decision-making processes in complex technological and societal contexts. More generally, workshop participants reiterated the need to articulate the ethical dimension *in* foresight activities (asking questions about the principles and values guiding future thinking) and *of* foresight activities (how foresight activities should be done in an ethically sensitive way).

We would like to thank the Dubai Future Foundation for welcoming us to the 2024 Forum Learning Day and to our workshop participants for joining us in this experiment. The feedback you have given us and the insights we have collected have provided us with significant learnings, which we will be incorporating into further iterations of this workshop design.

Annex 1. Traditional scenarios (group one)

First scenario – Genomic health prediction is used in special cases and under clinical supervision

It is the year 2030. Data protection regulations are strict and budgets for healthcare are high.

As a result of high regulatory and governance standards, and broadly well-functioning healthcare, most people are comfortable with the use of personal data and AI in healthcare contexts.

Health services use *genomic health prediction* sparingly, and under strict clinical supervision, to improve treatment and outcomes for the seriously ill or those who need to take particular precautions regarding their health.

Genomic prediction is used, where recommended by a clinician, to help people understand their disease risk profile (and how it might best be managed), and to make predictions about how they might respond to drugs or treatments.

Health services do not use *genomic health prediction* to try to reduce overall demand for services. It is not offered to the whole population, or to help people better understand and cope with their disease risk profile. Along with scepticism about the efficacy of 'personalised public health', health decision makers are not comfortable with deploying *genomic health prediction* at a scale that would make clinical oversight (and support for the subjects of the analysis) impractical.

Second scenario – Genomic health prediction is used widely for mass prevention

It is the year 2030. There has been a successful deregulatory push around data and AI, especially in the context of healthcare. At the same time, budgets for healthcare provision have contracted, leading to health providers looking for ways to cut costs and reduce demand.

Low levels of regulation and the increasing use of diverse kinds of data by insurers have made the public wary about sharing sensitive data for healthcare purposes. Many people are reluctant to share personal and health data for fear that insurers might use it to raise insurance premiums for some people -- or to restrict access to insurance altogether.

Despite these low levels of trust, *genomic health prediction* is available and targeted at the majority of the population – not just those who are ill or at high risk of serious health problems.

Genomic health prediction has been bought in to more efficiently allocate smaller real-terms budgets and preserve the time of an overstretched workforce.

The decision to pay for a prescription is often guided by pharmacogenomic insight, with drugs that show considerable genomic variation in efficacy prescribed only to those who it is predicted will show a sufficiently positive response.

Genomic health prediction is also used to enable far more targeted public health interventions than would otherwise be possible. Members of the public who share their genomic data are given personalised information about their genomically determined disease risk profile, so they can make personal lifestyle and other behavioural changes to minimise risk and nudges can be sent to their phones.

To cope with a shrinking workforce, the health services have invested in AI chatbots to replace the role of most medical receptionists and some diagnostic and therapeutic interventions, and as gateways to more specialised care. The data from wearables and genomic screening is used to compensate for the loss of rich contextual information gained from face-to-face contact with a medical professional. These AI systems can look at people's wearable data, self-reported symptoms and *genomic health prediction* generated predictions to make reasonably accurate assessments about where best to refer people or what drugs or medicines to prescribe. In cases of low certainty, these

systems they may refer patients to a human clinician, but many referrals and prescriptions can be generated directly through the healthcare providers' apps.

Insurance companies also use genomic health predictions to determine the price of and access to health insurance. Most health insurance packages require people to share their genomic data, to determine their insurance risks and therefore the cost and terms of their insurance. It is common for private healthcare providers to make insurance coverage for those deemed to be poor genomic health risks contingent on lifestyle changes to mitigate such risks.

Annex 2. Ethical techno-moral scenarios (group two)

[Scenario - PART 1]

2026

Scientific advancements have led to genomic health prediction (GHP) being used under clinical supervision, to improve treatment and outcomes for people who are seriously ill and those who have been identified as high-risk for particular diseases.

2028

Genomic health prediction is used to predict how people will respond to particular drugs given their DNA, with this insight used to inform decisions about the best medications to prescribe to people. A person unexpectedly dies from an adverse reaction to a type 2 diabetes drug that was recommended for them by a genomic health prediction system. It is revealed that the drug prescribed was not appropriate for this person. The incident results in a public outcry about the safety of these tests and the national regulator is put under scrutiny.

A patient charity says: *"They are putting innovation ahead of people. Every life should count".*

The CEO of the national regulator says: *"We are horrified by this incident. The introduction of new innovations always carries a risk, but if the risk is so low to justify the benefits for the many, then we have the duty to take it. Patient safety is paramount and we do every check possible to minimise these risks, but we also need to account for what a good for the population. We have opened an investigation and will be seeking to identify appropriate levels of compensation for the individuals and families affected in this terrible accident".*

An independent investigation rules that the risk of inappropriate drug prescription is sufficiently low and that the rollout of GHP should continue. However, some additional governance is recommended:

1) Patients may opt out of using GHP.

2) If a patient accepts GHP, the therapeutic recommendations it offers should be closely checked by clinicians – this will provide extra protections for the patient and reduce the risk of deskilling health professionals through the reliance on GHP.

3) There should be significant investment in additional training of clinicians. e. This will ensure clinicians understand how GHP works and can speak to their patients about its results.

2029

A study of clinicians using GHP reveals that half of them feel their training is not sufficient to equip them with the knowledge they need to critically assess its therapeutic recommendations.

In a letter to a national newspaper, a group of esteemed clinicians write: *"We are trained to care for and treat people, not to assess complex algorithms. With only 20 minutes per appointment we are expected to provide compassionate care while* also engaging *with and translating complex computer-based assessments. This should not be our responsibility. It is stopping us from caring well for our patients."*

To address these concerns, a national body develops guidance to assist clinicians in how they can incorporate GHP insight into their decision making. A key component of this guidance is for junior healthcare professionals to consult with one senior member of staff who will has completed relevant GHP training.

2030

A senior clinician caring for a patient with an autoimmune disease that is not covered by the clinical guidelines, deviates from the GHP recommendation and provides a different treatment. The patient suffers complications and sues the clinician.

A debate in Government is called and several arguments are presented.

The clinician membership groups claim the way GHP works means clinicians have little ability to interrogate or understand its recommendations and that means these computer systems are competing with clinical expertise. They worry that if the clinician is judged to be guilty it would lead to clinicians deferring more decisions to GHP rather than risk litigation when their judgement suggests GHP recommendations are not best for their patients. This will mean medical decisions will be increasingly subject to automation bias.

Research from a leading GHP safety Institute argues that removing clinician assessments from the therapeutic decision-making process is the best course of action. According to them, the human element injects more subjectivity into the system then is desirable. In their view, it would be better to only trust automated systems, narrowing the scope of the evidence that the medical system should routinely consider.

[Scenario - PART 2]

2031

Continued advances in advances GHP have created an emerging belief that it could enable the national healthcare system to become much more preventative, automated and streamlined. This narrative supported by pharmaceutical companies developing the GHP innovations that have been included in the healthcare system strategic vision for 2035-2045.

In this vision, ‘Every citizen would have a Personal Health Account that they control. It will store their health data, including self-testing and diagnosis data from whole-genome sequencing identifying known risk factors including a family history of disease.’

In order for this vision to become a reality, the majority of the population will need to have their genome sequenced and accept the use of GHP in their healthcare.

2032

Further work on the vision shows the national roll-out of it would require significant investment to update infrastructure. In response to this, the healthcare system evaluates a phased roll out where private insurance companies would offer their customers a chance to have their genome sequenced for free so that they can gain access personalised lifestyle advice. According to these plans, this information could be shared with public healthcare providers.

The healthcare system also explores opportunities for individuals’ genomic risk scores to be used in public health settings. The aim is for this data to help with triaging and decisions about who to prioritise for screening and diagnostic tests. According to these plans, this data could be transferred so that, with patient consent, any team working within the healthcare system or with those from a private provider can gain access.

2033

These plans, published last year, receive a negative reaction from the Nuffield Council on Bioethics (NCOB) and the Ada Lovelace Institute. They reject the idea that insurance companies should offer

this service as this would give them excessive powers over public health considerations. They say tighter regulation and governance is needed.

The Government, which subscribes to the vision say the use of additional regulation would be too costly and cumbersome. They agree insurer tests must meet high standards of quality, and that individuals enrolling in these programmes must be well informed about their proper use and limitations. But, they say that involving insurers is the best way of bringing GHP technology to the masses as it keeps national costs down and that in the end, the decision about how genetic information is collected and used is a matter for the individual to choose. In a public speech she explains: "*we must accelerate and adopt new advances in technology that can enable health professionals to make earlier and more effective diagnoses, alongside interventions that can empower individuals to take greater personal control of and responsibility for their own health."*

The plan to use private insurers goes ahead.

2034

A year after two insurance companies roll out the programme, a patient with a high-risk lung cancer score is denied financial reimbursement for asthma medication because it is revealed he did not quit smoking despite the lifestyle advice offered by GHP. This case receives media attention.

A spokesperson of the insurance company releases a statement saying: *"Individuals are in control of their health; they are provided with all the information they need to conduct a healthy lifestyle. With information comes responsibility. If they knowingly decide not to take up responsibility for their health, they need to bear the consequences for this."*

Public comments on the article appear to agree with some saying: *"Will this person now receive treatment through the NHS? This is not fair for taxpayers! Those who get ill despite being equipped with insight into their disease risk – having failed to act on this knowledge – are less entitled to taxpayer-funded healthcare".*

A campaigning association for social justice starts lobbying for regulation that will stop insurers making these decisions. They say it overlooks the impacts of social determinants of health, meaning that people from lower socio-economic backgrounds will inevitably have lower chance to conduct healthier lifestyles and should not be considered responsible and punished because of this. As they put it: "Being healthy is not always a choice!"

Human rights associations support the call, saying: *"Taking care of one's health, and being a good patient, may become a full-time duty even for those who consider themselves to be healthy. Not everyone can afford it and it will place extreme, unfair pressure on those with high disease risk scores to avoid burdening the NHS. Those with poor disease risk scores could be subject to much more onerous requirements to maintain access to insurance coverage than those who are lower-risk".*

As these campaigns spread on social media, an increasing amount of people become reluctant to share personal and health data for fear that insurers might use it to raise insurance premiums for some people -- or to restrict access to insurance altogether.