

## 2025 Royal Society of New South Wales and Learned Academies Forum: “AI: The Hope and the Hype”

### Panel Session 3: AI in Health<sup>1</sup>

Michael Blumenstein, Flinders University

Amandeep Hansra, GP and Australian Digital Health Agency

Simon Carlile, Google Research (Australia)

Richard Taggart, eHealth NSW

Sean Baldwin, NSW Health

#### Michael Blumenstein

I am the chair of what is likely to be one of the most controversial and exciting sessions of the day: the AI and health session.

What you might have noticed was a really positive vibe around AI leading up to this session. I’ve been researching AI for 25 years. And we’re in the summer, as Pascal said, but I’ve seen the winter of AI. We’re maturing in this country: when we talk about AI a decade ago, the then chief scientist, Alan Finkel, was quoted in saying that AI would be a threat to society, would bring down everything.<sup>2</sup> When you listen to predictions like that, back then AI was very different to what it is now. That that posture really created the cloud that has sat over Australia for a long time in relation to how it reacts to AI.

In an AI arms race across the globe, we need to be a contributor to the innovation, not posturing, waiting, and perhaps not doing anything. The danger of not doing anything will probably outweigh our worries and fears. It is really positive to see the change that has occurred: how people are

reacting and responding to the absolute explosion of AI in every part of our lives. People don’t know that AI is in their washing machines, in the brakes of their cars. It has been for decades because AI has been different over different periods of time. The future of X, as Pascal said this morning, is AI plus X. This is the reality: nothing is moving forward without AI.

It’s about how we respond to it and how we embrace it, rather than worrying what might be because someone isn’t treated as well as they could be in terms of their healthcare, or whatever. We’ve got to look at the risks for sure; we’ve got to look at governance; we’ve got to look at regulation, but we really need to update and evolve our posture, particularly in Australia. When you look at AI plus X — in this case, AI plus health — this is a really big issue, not just for the state, but also for the country and globally.

Recently, UTS commissioned a report a white paper with PwC,<sup>3</sup> where we basically did a survey of hundreds of clinicians,

<sup>1</sup> This is a transcript of a presentation, available at <https://www.youtube.com/watch?v=PYKhfcc0B7s&t=8s>

<sup>2</sup> Finkel A (2018) Chief Scientist warns AI could create “parallel society.” <https://www.chiefscientist.gov.au/2018/07/article-chief-scientist-warns-ai-could-create-parallel-society> [Ed.]

<sup>3</sup> “Reinventing healthcare. Unlocking the power of AI in Australia’s health and wellbeing sector,” UTS Price-WaterhouseCoopers, 2025. <https://www.pwc.com.au/health/health-matters/reinventing-healthcare.html> [Ed.]

healthcare workers, end users across Australia. Interestingly, what came out of it was there was apprehension, but there were also really interesting views from among the community around what the promises could be. What is really important about these discussions today is that we have experts who are sitting on our panels who are some of the best in their fields. They are also people who have crossed over between clinician and technologist, even expert in technology. It's really important, whenever you do AI plus X, that you have a really good understanding of what the X is, because there's no point in talking about AI if it's only from the perspective of a technologist or a computer scientist. It has to be the perspective of the end user, the actual adopter, who is going to make things work.

Today we have a panel that bridges that gap perfectly; everyone here has a background that is going to contribute to a very interesting discussion. I'd like to jump straight into a panel discussion with questions because I would like to leave at the end time for people to ask questions from the audience.

So without further ado, I'd like to introduce each of the panel members. Dr. Sean Baldwin is an emergency physician, Deputy Chief Medical Information Officer for North Sydney LHD at NSW Health. Dr. Simon Carlile is the senior staff research scientist at Google Research Australia, but is also in his own right an esteemed academic entrepreneur and will add a very interesting perspective to the conversation. Dr. Amandeep Hansra is a general practitioner, but also a chief clinical advisor in medicine at the Australian Digital Health Agency and has a significant background in bridging the gap between digital and

health. Richard Taggett is the Chief Executive eHealth NSW and Chief Information Officer of NSW Health. Richard will bring a very exciting perspective, as will Sean, around the NSW context.

This morning we had an AI and health panel and there was also (in the Law panel) a lot of attention around regulation and what that means. I would like to get this panel's perspective and start with Amandeep around the issue of balancing regulation and ensuring that is done in a way that doesn't stifle innovation, but with particular reference to health and healthcare.

### **Amandeep Hansra**

I've had a really interesting career journey and I've found myself at the intersection of exactly that problem. In my role with the Australian Digital Health Agency, even though we're not a regulatory body, we are responsible for the national digital health infrastructure. Our stakeholders are all the states and jurisdictions, as well as the Commonwealth government. We need to make sure that the foundations are right. It feels like we're building an amazing house, but I need to make sure that I'm laying the concrete properly so that all the systems are connected, and we're speaking the same language, that we can actually then build and innovate on top of that. But I also work as a GP. So I'm on the frontline seeing patients and seeing my peers start to adopt technology.

Another role is I work with startups. I work with a lot of very early-stage companies with founders who come up with great ideas of wanting to solve big problems and who are particularly attracted to healthcare. We help connect them with clinicians and then help fund them through the doctors'

network. We get great innovation happening, but the problem is that you can't innovate and scale that innovation if you don't have the foundations right. The foundations are not just the infrastructure that we're building at the agency — around connecting all of our systems, making sure we speak the same language, making sure that we can identify patients, clinicians, everyone across the whole system — but we are also making sure that we do it with a safety lens. When things go wrong in health, people die.

I'm not saying that any other industry is not important, but for us as clinicians, our main job is to do no harm. We have seen innovation move at rapid speed, with the regulatory bodies trying to play catch up. The problem is that if you do these things without proper regulation, you can end up with people coming to harm. A lot of the people building AI tools are very well intended, but exactly to your point, Michael, they haven't actually got the health-specific knowledge or they're not connected. They're technologists with great ideas — and I love when they come into the healthcare sector and try and solve problems for us — but they need to bring us in on the journey. So there needs to be a lot of co-design because we understand what the stakes are. We have some great work happening with the Therapeutic Goods Administration (TGA) which has obviously built some regulation around software as a medical device.

If AI has a therapeutic or a diagnostic role, then it must obviously be regulated. The problem is that we're seeing a lot of AI that is in a grey zone. For example, in my practice, I use an AI scribe, where I talk to the patient and then I don't have to look

at the screen and type notes. I can directly share eye contact and look at the patient, do my assessment, and the AI scribe has documented what's happened in that consultation. The problem is that sometimes when I look at the notes, they haven't directly translated what I've said, but have made some interpretation and said, "You've listed these symptoms. I think your brain is saying this is what you think the patient has." And that starts to get into diagnostics but these devices are not regulated.

We've got no clear line: what is safe and what is regulated. Many of our regulators are grappling with how they play in that area, but without stifling that innovation of those people who are coming into the sector, really well intended and wanting to solve the big problems that we have. It's a bit of a balancing act, but the key is making sure that you have clinicians involved in the design of these projects, of these innovations.

### Simon Carlile

I've got a question for you. One of the things that really impressed me about the AI and Law panel was that I came away with the notion that in some ways treating AI is different from everything else. It's not really the right lens. There are plenty of examples of where — if we're talking about discrimination or harm and those sorts of things — you can just strike AI out and use the legislation that we've got or rather insert AI in. There are still obvious complications around how you assign responsibility and those sorts of things. It brought to mind the late Daniel Dennett (1942–2024), the philosopher and neuroscientist,<sup>4</sup> who was asked about deep

---

<sup>4</sup> Dennett DC (2023) The problem with counterfeit people: Companies using AI to generate fake people are committing an immoral act of vandalism, and should be held liable. *The Atlantic*, 16 May. [Ed.]

fakes; he said, “But we have the mechanisms to deal with this.” And when probed, he said, “Well, look, legitimate currency is only legitimate because there’s no way of reproducing it. And people who attempt to reproduce it are hunted down and they’re prosecuted.” It’s an example of where we have a mechanism in place; it just needs to be repurposed. I’m wondering whether the regulatory things that you guys face in your own practice — because I’m a neuroscientist by training rather than a clinician — might be translated more rapidly than having these large open debates that go on interminably?

**Amandeep Hansra:** It’s a great question. When we’re building clinical governance frameworks, it’s the principles that matter, not so much the technology. If we spend five years writing clinical governance frameworks or accreditation guidelines, by the time we’ve written them, new technologies emerge and we have to rewrite them. We can’t write them about specific technology. My issue is: if you think about a drug that goes through a regulatory process with the TGA, they assess the drug, the ingredients, the manufacturing process, and then say, “This is safe for patients to use.” The drug that they’ve assessed is the same drug that goes out to the public. The problem is when you assess AI through a regulatory process, the AI is constantly changing. So once it’s out there, it’s not the same as what was assessed.

Sometimes that’s where we’ve fallen into traps: where there’s been an update or they’ve been dependent on open AI models that have changed, and then you realise that there’s a drop in quality or there’s a bad outcome and there are no real post-market surveillance processes in place. So we have to develop some new mechanisms because it

is unique, I think, to our traditional device and drug regulation.

### **Richard Taggart**

I would love to build on your question and comment on Amandeep’s. I’m really proud to work in the NSW health system alongside some 180,000 people. We provide care in over a thousand locations, from the city to the bush. If you’re a consumer coming into one of those services or you bring your family member into one of those services, you expect empathetic care that is kind, appropriate, evidence-based, safe. You expect that you get the best-quality care that you can. The way that we do that, the way that we ensure that we can be as good as we can be — and we’re not perfect — is through the governance, the regulation, the policies, the procedures, the structures, the frameworks, and the foundations that Amandeep just mentioned. AI has to coexist with some of these existing well-established foundations and build upon them and enhance them where it makes sense.

In fact, AI itself has to be a regulatory tool. I think that’s the opportunity: you wouldn’t expect as a patient that one of our clinicians would provide you a treatment that didn’t have evidence or didn’t have some level of security or safety around it. AI needs to have the same lens on it.

**Michael Blumenstein:** That is a really important point you raise about the evidence, Richard. When we’re talking about health, we’re talking about people’s lives. What’s at stake is significant. It’s different from someone just browsing on Netflix and getting the wrong recommendation by AI; the damage isn’t huge apart from being bored to tears for two hours. But we have a big issue if the AI is wrong and the impact is huge.

I'm going to touch on evidence in a second because I think that's a really important part of the discussion. Sean, what do you feel are the misconceptions that people may have in relation to AI-enabled healthcare that sometimes bring fears and concerns up in individuals who are pursuing the healthcare system with AI-enabled processes?

**Sean Baldwin**

Great question. Where I work in emergency departments, our patients come to us and they're unwell, so they're worried already. They walk in the front door, they're worried about what's wrong with them, about what's going to go wrong with their health. And will their clinician appropriately diagnose them? Will they treat them? Will the treatment have side effects? Is that going to cause them harm as well? And now they've got all these extra worries. They're worried about where their health data is going. Where's it stored? Is it secure? Is it going to be part of the next data leak reported in the media and who's got access to it and who can learn all about their private data? So they're worried on top of that and they're thinking the doctor might make a mistake.

What if the AI makes a mistake? What if it hallucinates during my presentation to the emergency department and now it's decided I need a brain biopsy? That's what they're really worried about. So your patient's worried, they're scared, they're unwell, and trying to explain to them, "Oh, but I'm going to use this AI tool." In that situation, having them understand what that AI tool does, how it integrates into the healthcare system, and how it affects them personally, is really difficult in a short time.

The AI scribes are a great example. All clinicians love them because they return us to

the bedside. They reduce our administrative burden, they give us more time to deliver healthcare and do the jobs that we love. But when this makes it into the media, all the commentary that you see back is patients saying, "Oh, I don't want the AI making decisions for me. I want my doctor to make those decisions for me." But that's not what the AI scribes are doing at all, although, as you say, the boundary is blurry.

**Simon Carlile:** This is something that really gets me going. I think this is a case of poor cognitive hygiene. The late Danny Kahneman (1934–2024) got the Nobel Prize in 2002 for, amongst other things, identifying what a range of what he call cognitive biases, which are ways that we think about things that are in daily life actually very effective. One of the things that he identified was that the ease with which you can bring something to mind equals your spontaneous assessment of the probability of that happening. That's when I stopped watching CNN because I renamed it the "Crisis News Network." This is about taking the narrative and making sure the sorts of things that we think are dreadfully important are really what are front and centre. That's what gets repetitive. Poor cognitive hygiene; we need to wash our hands of CNN, for sure.

**Michael Blumenstein:** I would say that news media and other sources of information have inflamed things. In the case of AI, there's a crisis around the corner or that there is a potential for real concern. And yet, when I talk to people like the esteemed people on this panel, I'm actually quite confident that there are a lot of things in place — whether it's governance or whether it's the ethical requirements in health and medicine.

I want to go back to the point of evidence. Would you both, Richard and Sean, and

maybe Amandeep as well, see a base emerging of evidence that there is improvement to be had when you use AI in the sector? And could you extend to the data that's used? Do you have some examples or would you like to comment on that?

**Richard Taggart:** I should caveat and say that I'm a proud public servant and work for NSW government. One of the things that the NSW government has done in their wisdom — because this is an emerging area — it has added new levels of bureaucracy, and nominated me as the person in NSW Health who has to fulfil that level of bureaucracy on behalf of the system, which means I often have to stand in front of clinicians like my good colleague Dr. Sean here and say, "Hang on just a second, show me the evidence of what you're about to do before we allow use in the system." That means that I've got a good view of what's going on out there across NSW Health. There are many hundreds of really quite exciting cases, but I would have to say the quality of the evidence that's showing the ones with real meaningful clinical impact is still at the beginning of the journey.

There is some really good solid evidence where things have actually gone through to fully regulated devices — there are some really great radiology use cases, for example — but the vast majority are still in that early phases. The projects I'm most excited about are where NSW Health clinician researchers are contributing to that body of evidence. But the opposite side of that coin is it isn't all just about healthcare and clinical outcomes. If I look at what's actually going into production, more than 50% of these use cases are in a category that I would call increasing the time for care. So, taking away the admin burden — we spend

almost a quarter of our time in the health system just on administration, keeping that thing running — anything that we can do to reduce that percentage of time — every minute that we give back to a clinician to provide frontline services — is a benefit.

There is good evidence that the time for care benefits for AI is pretty strong. There's some strong evidence in some of the administrative tasks, the finance functions and submitting forms and getting the quality of the submissions back to the various government agencies. That's why you're seeing very rapid adoption of those perhaps less attractive, but equally as important, use cases.

**Sean Baldwin:** I can definitely empathise with these discussions. I have them on a smaller scale — just within a health district rather than throughout the whole state, but all speaking the same language. I think that's the real problem I see with AI. Any treatment that I'm going to give to a patient is going to be supported by high-quality scientific evidence, but the high-quality scientific evidence for most of the AI programs that we see is lacking. There are a couple, but they're few and far between. When I see a large-scale randomised multi-centre trial evaluated using AI and evaluated not using AI and see that there is a patient-oriented outcome difference — which is morbidity or mortality to do with that treatment — then I'll adopt it and I'll use it. I think we'll see medicine shift that way as the evidence arrives.

I have no doubt that it's coming. The problem's going to be keeping pace with all the evidence, but I'm probably going to use AI to keep me up to speed.

**Amandeep Hansra:** I think that we also need to look again at the evidence in this new evolving world. Historically, we have relied

on the multi-year randomised control trials. The problem is that, for example, from the moment when you identify the actual target of a drug to going into production, it's something crazy like 17 years. We can't do that with AI because by the time we do some of these trials and we publish them, the world has moved on and we're already onto ChatGPT 1500. We're moving so fast. I'm not saying that we shouldn't have evidence. I think it's just that we need to look again as an industry at what is acceptable evidence in this world where we are moving at such rapid speed. Are we comfortable with that? How could we use AI to help accelerate how we're generating some of that evidence?

The other thing that I'm finding is that we're getting a lot more real-world evidence. When you think about the evidence that we have from some of the trials — for example, we haven't had women represented in a lot of research trials over the years. It has only been more recently that we've really pushed to have gender equity in terms of representation. The way that I was taught to pick up a patient having a heart attack was based on mostly men's symptoms but women present differently.

We haven't always had cultural diversity, First Nations representation. But what we're finding, with the world of AI and the way we're capturing data and how we are able to take some of the meaningful insights from that data, we're actually finding more representation of our communities in the data. So we're actually able to see what that drug or that intervention looks like in the real world because I'm actually now checking what it's doing to my heart rate. I'm checking what it's doing to my sleep. I'm getting a lot more patients who are coming back to me with actual data of how that intervention affects

them individually because they've got access to technology that they didn't have before.

To me, that is actually much more important than what randomised control trials have shown in some ways because it's personalised care. It's about that individual and how it's impacting them — not a whole group that wasn't representative of the patient in front of me. I'm probably going to get shot down by a lot of my research colleagues who say, "What? How dare you say this?" but I do think we need to stop, and take a moment and ask, "What are we able to accept as evidence moving forward?" Health is the second largest industry in terms of publications in AI outside of technology. So we're doing great in publishing a lot of things about AI trials, but the problem is we're not also scaling those. We've done the trial, we've got a bit of evidence, but then how do we take the next step and actually put it into the real world and say, "Are we feeling comfortable enough to trial this on a state basis and a country basis and a large population?"

**Richard Taggart:** I think that's just an Australian problem. We're very good at producing the evidence, but very bad at actually adopting it.

**Michael Blumenstein:** I really think that the perspectives here are so exciting to hear, given the fact that there's a representation of someone who can oversee what's happening in NSW Health around technology and may making sure that it's going to be for the betterment of patients in healthcare. I really want to emphasise what you said, Richard, about the backend, the non-direct-patient-affecting AI processes that are occurring. The single largest expenditure in the NSW budget is Health. So there are opportunities around, and there is actually a lot of

advantages outside of just direct health and patient care around advancing the operations of facilities, of hospitals, etc.

The other thing is Amandeep's point around the question of what we are looking at in terms of evidence. What is the data telling us? And now there are so many places for data to come from. I want to ask Simon: how can we scale up what we're doing? What are your thoughts around the role of universities, the tech industry, healthcare providers, and governments to contribute to that scaling? What else is there that brings together important institutions to drive this innovation, that can contribute to the future?

**Simon Carlile:** I'm very happy to speak to that. I wouldn't want to lose the previous track, though, so we can come back to that because I think there are some real changes afoot that change the focus from illness to wellness. That is going to happen. AI is going to play a big role there, so I'm happy to chat about that later. But, okay, returning to universities. I've been a university professor. I've worked in the area of hearing health and researched in that area for most of my life. I've had spinoff companies and I'm now working for a large technology company, so I've got a little bit of visibility on each of those things. I think that universities in particular — and it was wonderful to hear about EduChat or ChatEdu or whatever they decided it is today, because that is an example of where we need to drive continual pedagogic reform in the way that we approach education, and no more so than in universities.

I had the pleasure of working with a number of brilliant clinicians and educators in the '90s where we designed a medical course that was based on problem-based

learning. The questions that we drew out were, how do you do differential diagnosis? How do you know what the questions are that you need to ask? Education should be about the process of learning and discovery, not the regurgitation of somebody else's facts, because by the time you get out to practice, they're irrelevant anyway, because the world has moved on to pedagogic reform that really focuses on the way that students and professionals can leverage AI when they have left the university.

In terms of university research, they're making an effort in this regard, but they're still very discipline-bound. As you would know — I was a pro vice chancellor as well — there are the baronies and the deans are the barons and there's a lot of fragmentation or silo-ification in universities. And while there are attempts to bring together much more collaborative in that direction, I think that can go a lot further. There are also Commonwealth incentives that could be put into that — not just dropping a little bit more into the National Health and Medical Research Council Fund for a so-called collaborative research project. I think the most salient thing that struck me in the years that I've been in technology is they've got to use a lot more user-centred design, particularly engineering companies.

I'm a scientist and a time-zone engineer as well, and I fall in love with my own solution, but is it the right solution? Is it actually moving the needle on the thing that counts? Anything that's involving the end user is super critical to make sure that your effort isn't being wasted. I think tech companies have got a lot to learn there.

I think health providers as well. I just came back from seven years in the United States, and that was an interesting experi-

ence for a whole variety of reasons, trying to navigate the health system there. What really brought into very sharp focus for me was the way one needs to really balance patient-centred care and — obviously in the US health system — a profit motive. And it's got to be win-win, not zero-sum. There really has to be thought about that. There are increases of efficiency but that shouldn't just hit the bottom line. How do we turn that back into increasing and improving patient care? I'm quite sure that while we're not purely profit-driven, we are budget-constrained. Many of the same arguments pertain at state and commonwealth levels here in this country. So I think these are still very pertinent arguments.

Coming back to regulation, it was raised this morning by an ex-minister. Slow is a problem. About three weeks ago, I had an internal briefing in my company about hallucination — we don't use the word "hallucination" anymore, not because we don't like the word, we call it factuality, because there is a huge amount of research that is going on, on determining whether a statement from a large language model or any kind of foundation model has a factual basis. We heard one example: for instance, we ask another LLM, is that true? Designs now are moving towards looking at this in a parallel process, so you don't wait for the answer and then ask someone else, "Is that answer a good answer?" It's all built into the process. So factuality is one thing.

The other thing that was raised today is that you don't understand what's going on inside this black box. There's no way of looking at the logic. Well, reasoning is also a hugely, hugely influential area in foundation model design and implementation now. They're not there yet, but to

your point about the way things move so quickly — and that's another real difficulty as well — is at what point do we start making judgements on the data that we have. Because we know that I'm not going to wait for the next super-fast CPU to come out before I buy another computer. I need a new computer now, and so I've got to make some risk assessment of where I am. They are just some quick thoughts on the four areas, universities, technology, health providers, and government.

**Amandeep Hansra:** Simon, I just wanted to comment on one of the points that you raised. I'm going to mention the word "productivity," sorry. One of the things that happened when we were being sold AI scribes is that the tech companies would come and pitch to us and say, "You can see this many more patients if you use an AI scribe." Obviously that will translate to billings for us in general practice. They thought that's what sold us and that's what made us pay the quite steep prices initially. The thing is that we didn't see more patients. We love the scribes, but we didn't fill that spare time with seeing more patients. We actually just spent more time with our patients. Sometimes we measure the wrong things — when you couldn't spit out the data and say, "This is how much money you're making extra or how much extra productivity." What happened was that the doctors went home happy, the patients went home happy, and the doctors wanted to come back to work the next day.

That is a measure of productivity, though, because we have a lot of doctors and other healthcare providers in the system who get burnt out and leave. And we've got about a 10% attrition every year of medical doctors who leave because they just don't enjoy

their work. Sometimes we think that we're addressing one thing, but, if you actually look at it, yes, there has been a productivity benefit, but it's not where we thought it was going to be. It's reducing that what they call in the US "pajama time," having people be happier at work and actually provide a better service to their patients and ultimately have better patient outcomes, which also leads back to productivity. I think we need to think more broadly about the benefits of AI rather than this direct return on investment (ROI).

**Richard Taggart:** Amandeep, you just touched on something that's at the top of my problem list at the moment, which is AI scribes. And if you're not familiar with the term "scribes" — effectively you allow the AI to listen to your consultation, then it writes up the note in a format that you want. So you can put it straight in your clinical information system and do various tasks with it. Actually they're becoming more of an agent now than they were. They're doing other tasks and there's a lot of very rapid adoption. This technology has been adopted by general practice, by consultants in their private rooms, by allied health professionals all over the world. The demand has been sensational. It has been really quite exciting to see that, but the implementation of a piece of software or a new tool or a workflow change inside a general practice, even a large practice, compared to implementing something in an enterprise as complex as NSW Health, are two very different scale problems.

There has literally not been one single day since I started my role on the 20<sup>th</sup> of January that I haven't had a request for a scribe of some flavour to be rolled out at scale in the NSW health system. I'm very eager to do it

because I can hear demand from clinicians a potential benefit, but it has many layers of complexity in the scaling of that technology and in implementing it safely. I'm sure this is something that comes up in your world as well, Sean.

**Sean Baldwin:** Absolutely. I'll send you an email later.

**Michael Blumenstein:** You guys just sit next to each other, pass a note, no AI. I would have liked to keep this going because you probably agree with me that this is a very interesting area: AI and health, because health touches everyone, no matter what age you are, no matter what size of family you have or how old your parents are.

The last question I'm going to ask (and it's terrible for me to do this). For the sake of productivity, I'm going to ask each panel member to give just a one-sentence response to this: someone mentioned the fear in the context of jobs with the rise of AI, in the context of health and the health sector. What do you each see will happen to jobs in the future with the adoption of AI?

**Richard Taggart:** Augmentation is the word I would use. We are not going to be able to meet the demand for the health system over the next 5 to 10 years. We've got a rapidly ageing population, which is wonderful and that we're living longer and hopefully healthier lives, but that increases the demand on the system — there aren't enough workers whom we can bring through the system, there aren't enough hospitals that we can build to satisfy that demand. We have to embrace AI so that we can get some level of productivity — to use that word — to meet the demands so people get high-quality care at that scale.

**Sean Baldwin:** Completely agree. I guess we work for the same organisation. I don't know anyone in the public health sector who's worried about losing their job to AI. I think if anything, we need to get more done with what we've got, the resources we have, and hopefully AI can help us do that.

**Amandeep Hansra:** Yes, I definitely don't think I'm losing my job anytime soon. I think I'm actually going to go back to providing care as opposed to doing lots of admin tasks. I'd love for AI to do all the stuff in the background that needs to be done so I can just sit there and focus on the patient in front of me, provide the human touch, provide the care coordination, be their partner in their journey through the healthcare system. I think AI is going to give me the opportunity to do that. So I'm hoping we can all fall back in love with our jobs.

**Simon Carlile:** I think that the biggest impact of AI is going to have nothing to do with jobs. I think it's going to have to do with this shift from illness to wellness, as I mentioned before. Everyone who's wearing a digital watch probably has a range of physiological sensors on that watch. You've probably also ticked a box when you put it on, which said, yes, I want to use the Fitbit notification of whether my heart rate's going up or down or blah, blah, blah, blah. So you've just given away all your data. (Well, not given it away because there are very strict privacy rules around that as well.)

The agents that are being developed by the companies that make those watches are looking at that data are not providing you a medical diagnosis, but they are providing you advice. They're also providing you indicators when you should go and seek some medical diagnosis. What they're building is these large models that can look at your data

and determine when it is starting to deviate from your standard norm and starting to indicate that you do need to talk to a medic about how you manage your behaviour or your diet, or whatever that's driving you out of the norm.

That's not going to impact anybody's job other than maybe decrease the load of chronic disease that is hitting our medical systems, increasing the rate at which people are able to provide diagnoses and treatment, because early treatment is always a much better way than waiting till it's stage four. So, yes, that'll be the biggest impact. It's happening now and there's no government regulation standing in the way and people are happily signing over their data and taking the advice of those agents.

**Sean Baldwin:** Certainly I see someone every day whose Fitbit or Apple Watch has told them to come to the emergency department because they might be dying. It's extremely rare that I find much wrong with these people. It does happen. Most of the time it's a measurement error or something's gone wrong, but it has made them so, again, anxious and stressed about their health because it's such a constant thing that they're getting exposed to every day. They're going, "Oh my goodness, what could this be?" And then they come in to be reassured by me.

**Michael Blumenstein:** I think that the formal panel session has ended in a very positive note, before we go to Q&A to the audience. I think we have time for a few questions.

**Q1, Susan Pond:** My question is to Simon, but it may overflow to the others. If you're at an industry conference — not one that's a more for-profit-driven conference — and you are a part of the research team now in Google, what are the big data companies

(if I can call Google that) thinking about in terms of making profits from AI in health-care?

**Simon Carlile:** That's a really interesting question. If I knew the answer, I'd probably be a vice president. My own research is actually applying machine learning to re-imagining and computational modelling to re-imagining the hearing aid. I'm a neuroscientist and a hearing scientist by training, so that's unfortunately what my focus is. So I'm not sure that I can give you a cogent answer from an industry perspective.

**Q2, Mark Pesci:** There's a new device that's just come out from a company called Kohler in America which manufactures most of the toilets that Americans use. It's called the Dekoda.<sup>5</sup> It's a camera that points down into your toilet tank. So when you've peed or had a number two, it will do an analysis. It uses edge AI and uploads it to the cloud. There's now going to be a galaxy of devices like this that will be flooding all of our medical systems with AI data around patients' wellness. How are we going to be able to incorporate that into our medical systems workflows?

**Richard Taggart:** I might try and tackle that question, although I'm fascinated by the experience. You must hold your breath when you use that restroom. I think there's a lot of opportunity here and a lot of upside. One of the things that I think about as a leader in the public system is how could we possibly keep large numbers of people away from seeing the lovely Sean and potentially have a regular appointment with the lovely Dr. Amandeep instead? There's one risk factor that I think would really move the needle on that scale — we could use AI devices and technology — and that is hypertension.

About 30% of people have hypertension, so I divide this room into three. Maybe you look particularly healthy, maybe you've got hypertension, and about half of you would not know. But if we could tell the half that doesn't know that actually you might be at risk, go and get your blood pressure checked by a health professional, we could potentially prevent a whole heap of downstream health issues. That's a relatively simple use case, but if we can scale that, I'm very happy.

**Amandeep Hansra:** I went to that health conference — I think it was in Vegas — and I saw that toilet. I was too scared to sit on it. But I think we're in this period where we've got too much data. People are being given all of this extra information and they don't know what to do with it. The clinicians don't know what to do with it either. I get patients coming in all the time with, "Here is my output of all my wearables, my Aura Ring, my watch data, spreadsheets of information." People who are not diabetics are wearing continuous glucose monitors and saying, "Hey, I went for a walk and my sugar spiked. What's wrong with me? I must have diabetes." I'm like, "That's normal." So the part that's missing is actually the actionable insights and getting those accurately.

We're going into this prevention space where we're going to arm people with their own data; with that, we have to educate them. We have to tell them to help them figure out what is something to be alarmed about that needs to send you to the ED or send you to your GP. I think people should be allowed to manage themselves. That is what I'm hoping as we start to give them access to their own data. One of the things that happened a couple of weeks ago

---

<sup>5</sup> See <https://www.kohlerhealth.com/dekoda/> [Ed.]

is we just changed the legislation for the MyHealth record — I hope everyone in this room has a MyHealth record — where we've now implemented "share by default," which means when you go and get a blood test, you should be able to access that pathology test result the same time as a doctor accesses it. If you get your blood done in the morning, by the afternoon you should be able to see that report yourself. A lot of people panicked about this and said, "Oh my God, people looking at their own blood test results, God forbid they're going to get panicked and they're going to end up in the ED." I thought this was a very paternalistic view — only I can hold your data, even though it's about you, only I'm allowed to look at it and I'm only going to give it to you when you sit down with me and I can bill Medicare and that we are going to work through it together. In some situations if it's complex, I agree, but I think we need to move away from being the custodians of all of our patients' information and giving it back to them, and educating them in the process.

We're going to go through a period where there's going to be a bit more utilisation of the health system before we get to this period where people are going to start to take responsibility, accountability, and learn about their own body in ways that they've probably never known before. So I'm excited about that toilet — a little bit scared, but also excited.

**Q3:** How important is context to diagnosis? At times you're saying that all the data being collected is a bit too much or it spooks the patient, etc., but how important is it to the clinician to get the right amount of data? I did a little bit of research on snoring. Why does somebody snore and what kinds of

snoring are there? I'm just wondering, we found it important to have something on the person to collect all the snoring patterns and the condition of the environment. That was quite a useful tool for the clinician to decide there is a problem not with the person, but in the way he was living. Same thing happens in aged-care centres, the CO<sub>2</sub> levels just go way up because all the doors are closed.

**Amandeep Hansra:** Look, I think that is part of our role in this world, where we're going to have a lot of data, where I say, "I'm not going to go out of a job." It's because that is some of the nuance that I'm going to have to figure out as a clinician what the context is. When was that information collected? You had a terrible sleep, but I knew that you'd been delivered some bad news that day at work. Maybe the AI doesn't know that, and maybe that is my job, understanding the human context, also understanding what's happening in your social environment with your family, your kids. This is why I think that we as clinicians are going to go back to doing what we're doing better, because we're not going to be gathering objective information that can be gathered by AI and wearables and systems; instead, we're going to be gathering the information that is more about the human and the context.

Going to your point about the environment, that is also static information that can be measured. We're now starting to look at smart homes, particularly for aged care, where I've seen people develop mats where you stand on the mat every day to brush your teeth and over time it can tell whether you're starting to develop osteoarthritis because you're starting to not have equal balance on both feet. We're going to start to see these environmental data collectors that

are going to feed into what the clinician gets to see in terms of your home living conditions, your context. How often do you open the fridge? How often do you move in the house? The problem is that at the moment we're not set up to receive all that data and interpret it. At some point, I would love it

if the AI systems could pull that together and give me a really succinct summary of "this is what I think is going on," but then contextualise it with the human factors that AI can't figure out. It's going to be a fine balance, I think, getting there.

