Tech Policy Exchange 2023: Keynote Speech

By Gregory Francis, CEO, Access Partnership

In case you are new to Access Partnership, our mission is to make sure as many people as possible have access to transformative technologies without impugning competition or consumer choice. It's a narrow but fun path we've laid for ourselves, and we use policy as our toolbox.

Because we do that all day every day, we know that even technology that creates the most positive transformation is pretty much as nothing unless it's married to policy that can help it take root and ensure users trust it. If that doesn't happen, the outcome isn't complicated and we don't need to over-intellectualise it. Without a marriage of good technology and effective policy, capital is invested unevenly across geographies, opportunity is shared unfairly, economic lacunas open up between people and societies, and the old rules apply: of might and right and wealth and power prevailing over fairness and shared human advancement.

We spend a lot of time in our company predicting outcomes and then walking back those outcomes to see if we can get to something better so that, at worst, innovations are appreciated and, at best, they are improved and form the basis of advancement and achievement everywhere.

But like all of you – and I am sure of this even if I don't know everyone – we have dragged our feet on preparing the world for AI. Even if we know what happens now, we also know how to improve the outcome. You all know what happens, too, I suspect, but shall we paint a picture?

Start with the present and the race that's on among some countries to use AI to maximum advantage. What we can be sure of is that those countries, where AI was or is being developed, will get the most out of it. That one is easy.

Where AI was not invented, it will have to be imported, sometimes by stealth, sometimes straightforwardly. To do that, AI exporters will say it is safe and good and transformational, and they will mean it because that *can* be true, but only in so far as a scalpel is all those things to a surgeon, but is otherwise dangerous unless it's used with care.

And dangerous things *will* happen at some point, probably in a country that's not the US or the UK or anywhere in the G7, and then the export of AI will be reviewed and *curtailed* in lots of places. AI importers will be told to stop or disentangle AI from everything else, to *throw away the lesser part of it*. But this will not be possible any more than it was for Hamlet's heart.

Then, there will be howling from AI developers and users, and policymakers around the world will start to impose hasty, ham-fisted rules that try to retain productivity gains and ensure safety.

Various AI purveyors will trade on their USPs: Chinese, Saudi, and Indian AI will be targeted at countries of the G77 and distinguished from the raw American stuff that, after all,

wasn't trained on local data sets the way China's was. Then they get into a shouting match: Big Western Tech will claim a more principled use of data to train AI, boast about better guardrails, and ask consumers: whom do you trust more, those who work with the avuncular Joe Biden or the denizens of Zhongnanhai?

Over the medium term, countries that do not ban or only lightly curtail AI will get the benefits but also suffer the most traumas. And like a population reaching herd immunity, there will be dislocation of whole workforces and challenges to civil liberty, but also miraculous inventions and solutions that, in simple terms, place the AI-using countries ahead of others.

At which point, unwilling to suffer the growth imbalance any longer, most markets will fold. And like so many lawn chairs, they'll accept AI again, suffering the benefits and the blows, and, for lack of choice, live the AI boom-and-bust that the transformation will wreak on their economies – only they will suffer it later than us, and bear it slower, and from a place of less bargaining power and, frankly, more resentment.

So our baseline task now, and part of what we are challenged to cover today, is how to make sure that 10 years from now, even they, the last movers on AI, will grow stronger, slowly richer, and develop their own AI, maybe applying lessons learned from the AI-indigenous economies, such that over a 15-year period, a more definite version of the voluntary commitments articulated on the White House lawn this summer will be in place the world over.

Now, to get into the practice of making sure our work today is not too US- or UK-centric, let's run a quick experiment looking at a couple of those commitments made by US tech giants and consider how they will be perceived outside the US, UK, Canada, France, and so on. It's a quick-and-dirty way of testing how quickly third markets will accept AI from tech companies. To do this, I'll pretend I'm a stakeholder in Malaysia reading a couple of these fancy principles.

We commit to prioritising research on societal risks that AI can pose.

I, the Malaysian, hear this as how to create process efficiencies and help office workers in Manchester and San Francisco to stay relevant. Maybe its intent is to help them deal with diabetes. But where I am, big societal risks are child hunger and ethnic and confessional tension. And I have absolutely no confidence that these are the societal risks that the White House was talking about.

At worst, I'm imagining that AI is going to help Chinese trawlers optimise their fishing in my seas, depriving my population of protein. I'm imagining how fast my country's Muftis will demand a ban on AI because it lacks a societal framework that's in any way Islamic.

Let's take one more...

We commit to publicly reporting AI systems' capabilities, limitations, and areas of appropriate and inappropriate use.

In Malaysia, we don't publicly report commercial problems, and I don't believe that anyone else does so voluntarily either. This sounds to me like virtue-washing. I'll never be getting the full picture on the risks of AI, not least because it's the AI company that will define what's a limitation and what isn't.

Anyway, what's inappropriate use? If it helps me identify fraud in my growing economy, then I'm going to use it, and if some or all of civil society gets caught in that dragnet, that's just the cost of development. I need to take the macro view, the same as China does, or I'll be crushed.

You get the idea. My purpose here today is not to deride any work that has been done to shape rulemaking around AI. On the contrary, at this point it is irresponsible to do anything other than seek to build on any reasonable foundation; there can be no pride of authorship at this time of urgency – just selfless contribution. So part of *our* contribution today has to be to understand how the rest of the world will see policies we are developing here in the UK, to make sure that at each step we decrease the chances of the inauspicious unfolding I described earlier, *and* in fact, hopefully go further: our work needs to enable the *dissemination* of AI across the global economy so that scenario I opened with does not come to pass. It's a modest proposal that, I know, since frankly, we should be calling for a Geneva Convention on AI, a whole Bretton Woods construct to deal with it, or something similar.

But we know the international community moves in baby steps. So, for today, and maybe for as long as we can hold the thought, remember that whatever we put in place should consider how others will read it so that we build trust universally and make it easier, not harder, to experiment with AI locally, the better for people *locally* to know how to optimise its uses for themselves.

If we can do that, if we can start to anticipate Malaysia's scepticism in our process of designing policies, we'll have earned our lunch and contributed to the work that this country is doing to put the UK at the centre of the fair and lasting development of AI and the policies that will give it reach.