

[DAISY THE GREAT:](#) [00:00](#) [Built my home on hollow ground.](#)

[CRAIG:](#) [00:07](#) This is [Craig Smith](#) with [Eye on AI](#), a podcast about artificial intelligence. There is a race underway between the world's great powers to integrate machine learning into their national security infrastructure and to develop artificial intelligence powered defense capabilities. This week, I talk to [Brendan McCord](#), who wrote the [Pentagon's AI strategy](#) and is now a Special Government Employee at the [National Security Commission on AI](#). Brendan talks about what he believes the US needs to do to stay competitive with China and promote an alternative vision of AI-powered security and prosperity to the world. I hope you find Brendan's ideas as thought provoking as I did.

[CRAIG:](#) [00:58](#) Okay. So, let's start with who you are, your education and how you got interested in AI for national security or national security as it relates to AI.

[BRENDAN:](#) [01:10](#) Well first, thank you for having me. I'm Brendan McCord. How I came to work at this intersection of AI and national security is that three or four years ago I sat in an audience and I heard then-Secretary [Ash Carter](#) talk about the Department of Defense being at an inflection point when it came to artificial intelligence. As I sat there – I was bringing with me a mix of engineering education at MIT and entrepreneurial education at Harvard Business School, experiences first in the operational military, having spent over 600 days underwater on submarine missions at the forward edge of the military, and then in leading a team of AI experts and software engineers at a startup trying to apply AI to public safety challenges– so, I heard Secretary Carter's statement as a clarion call for me to join the Department and work on that mission.

[BRENDAN:](#) [02:11](#) So I joined the [Defense Innovation Unit](#), based out here in Silicon Valley, as a technical appointee, as head of machine learning. I worked there on really interesting open source projects, built big datasets, pursued world speed records for training deep neural nets on public infrastructure. That gave way to solution building type things. I helped manage the Google effort within [Project Maven](#) and then also helped to apply an interesting non-deep learning, imperfect information, game-solving AI to a military planning challenge. Ultimately though, as momentum built, that role led to strategy development and to the formation of a new team in the Department of Defense called the [Joint AI Center](#). So, the end result, as I look back - and I'll say here that these are my personal views not those of the Department of Defense or the

government - but as I look back, it was a front row seat on history. To some degree I got to see this pivotal point of an AI focused transformation of one of the largest, most complex, and most entrenched bureaucracies in the world – but one with a crucially important mission.

- BRENDAN: [03:31](#) Today I'm at a company called [Tulco](#). I work with billionaire entrepreneur [Thomas Tull](#) and with a wonderful team of data scientists and machine learning experts. Our mission is to transform sleepy parts of the economy. We're an alternative investment platform structured as a holding company, and we work in sectors such as medical apparel and security, waste management and other areas. We partner with challenger companies led by bold entrepreneurs who want to go take these static sectors on and transform them, using advanced technologies and the new business models they enable. Our AI lab helps them do that through a partnership model, helps them accelerate, compete, and win.
- CRAIG: [04:14](#) Wow. Okay. That's quite a lot.
- MUSIC: [04:18](#)
- CRAIG: [04:25](#) So, you left DOD what year?
- BRENDAN: [04:28](#) Earlier this year.
- CRAIG: [04:29](#) I've seen you referred to as the architect of the DOD's AI Strategy. What does that mean and how did that come about?
- BRENDAN: [04:40](#) Well, to frame the strategy, I'll start with the historical lens. For the last five decades or so, DOD's focus on AI, and it was focused on AI, was on enabling risky discovery and on fostering technology creation with early stage research, basic research. We viewed AI as a research concern almost entirely. And the work that was done there was nothing short of visionary. If you think about the [Defense Advanced Research Projects agency, DARPA](#), their work in perception and natural language, planning and so forth, and the partnership they helped build among the government, academia, and industry led to personal assistants and prosthetics that are near natural, and self-driving cars. To say that that work transformed technological practice would be to really understate the case. It transformed the world. But I mentioned that we were at an inflection point.
- BRENDAN: [05:50](#) What we realized was that we had to now translate the technology, translate AI into decisions and impact and had to do that at the speed of technological advancement and at the scale

of the entire national security enterprise. So, we created a strategy to do that. I'll talk about three aspects. First, much of the insight from the DOD strategy is actually around the art of organization. For example, we had over 600 active projects in AI and each one of these projects was globally unique in the software sense, from how data was managed, to development work, to how things were fielded, to acquisition on every dimension. So, this is incredibly complex and it's slow and it's costly. And so, to go from that to an enterprise approach in which decentralized projects flourished but on top of a common foundation was one aim.

- BRENDAN: [06:59](#) Second, how do you deliver quickly? I believe you use the approach that has been followed here in Silicon Valley and in decades past in national security when it has mattered. You have a small team, you give them a specific problem to solve, you prioritize shipment date, you work directly with your end users, you start small, you field, you iterate, and then you scale until eventually the entire mission area is transformed. So, there's a lot of common-sense approach there. The third thing I'll mention from the strategy is around leading in military ethics and AI safety. One of the clearest statements that you'll see in the strategy is that as the Department moves out, its focus will include ethics and humanitarian considerations, and this notion of short- and long-term AI safety.
- CRAIG: [07:50](#) Yeah. And then out of that strategy came the JAIC. Is that right?
- BRENDAN: [07:55](#) That's right. The JAIC is the Department's new AI center of excellence. The mission is to transform the DoD by accelerating the delivery and adoption of AI to achieve mission impact at scale.
- CRAIG: [08:16](#) So part of the challenge was mapping all of the projects? Understand how they fit together into an overarching strategy and then create an entity that could coordinate among them? Is that right?
- BRENDAN: [08:31](#) Part of the challenge was, as you mentioned, to prioritize and create coherence. And one of the ways that we attacked prioritization was to name a small number of so-called '[national mission initiatives](#).' These are big problems, big mission areas that were ambitious in their scope. Examples include a national mission initiative for humanitarian assistance and disaster relief, a national mission initiative for predictive maintenance, to change the way we maintain our equipment wholesale in the Department of Defense. And rather than having small hobby projects, so to speak, these are large efforts of a certain

structure designed to both really change that mission area and to advance this common foundation, so that projects thereafter can flourish on top of it in a highly decentralized way. Without an initial wave of NMIs, it's difficult to put that model in place.

MUSIC: [09:30](#)

CRAIG: [09:36](#) there's been a lot of criticism in the press or by people that look at national strategies across the world, that the U.S. does not have a well-defined strategy in the way that China has, or Russia has, or even France has. Is that true? I was talking to [Trae Stephens and Brian Schimpf](#) and one of them made the point that, particularly China and Russia are structured - their economies and their political systems - in a very different way, so that having a centralized strategy makes sense. Whereas the U.S. is by nature decentralized. I'm just curious what you feel about that.

BRENDAN: [10:25](#) Well, I think that good progress has been made on the military and intelligence side of the equation. And while I think we are starting to solve for speed and solving for organizational change, which I do think were the principle challenges we faced, I think more needs to be done at the national level.

CRAIG: [10:47](#) In terms of funding or in terms of organization?

BRENDAN: [10:51](#) in terms of strategy, in terms of national strategy, and committing as a country to resolving and dealing with the opportunities and challenges that AI poses. And broadly speaking, I've heard those broken those into three. I would say there is one around military and intelligence effectiveness and the revolution that's occurring there. The second is around economic competitiveness, which includes things like technological job displacement. And the third is around the future of democracy, which includes many things, but certainly individual freedoms that are affected by AI systems that are fair, accountable, transparent, and ethical in operation. So those three different aims, are all things that have to be dealt with in the United States in the context of our western democracy. And so, I think much more needs to be done in that regard to deal with things, particularly outside of the military and intelligence arena.

BRENDAN: [11:52](#) Just to give you some examples, I think that if the United States were to proceed boldly there and make resource commitments for dealing with those challenges and make political commitments for dealing with those things, I would want to do a handful of things. One, I would want to attack the entry

barriers to progress, and this starts by focusing on increasing the supply of skilled AI professionals. The single most significant entry barrier in this country and around the world is the scarcity of talent. You could increase the number of U.S. computer science students by a factor of three or five. You could also, with about a billion dollars a year of funding, I believe, incentivize our companies to build the workforce needed for the future and incentivize them to do worker training with a tax credit. You could better harness the gift of global talent with new pathways for high-skilled talent from China and elsewhere to live and work in the United States.

BRENDAN: [12:54](#)

And you could expand the discipline with efforts aimed at diversity and inclusion. So that's one area – human capital. Another type of entry barrier in AI is data and the physical infrastructure of AI. By scaling up access to infrastructure for innovation, you allow academia to flourish and you let new firms better compete. Just as an example here, imagine an American cloud credit where there's a free-market, credit-based approach to accessing those important resources. The government also has a huge role to play in both unlocking existing data with privacy-preserving AI, or interoperability, or resolving copyright issues and so forth, but also in provisioning new data sets that maybe would never have existed, tipping the scale on data that would never have existed in the world. Satellite imagery is a great example of this, which of course led to a revolution in mapping and many other things.

BRENDAN: [14:12](#)

To foster sharing of data, you could charge universities with acting as a network of data brokers. I think other important areas for the United States beyond the entry barriers of talent, data, and compute are what I think of as AI for X, and this is really inspired by some of the work that people like [Tom Kalil](#) have done or what they're doing at MIT with the [College of Computing](#). Examples are AI for scientific discovery, applying it to things like the material science or chemistry. Another, in a more practical action standpoint, AI for X could mean AI for global public goods, like humanitarian assistance and disaster relief, or for healthy longevity. Imagine if you had longitudinal healthcare data from DOD and from the VA and from Tricare. It would be incredible for that.

BRENDAN: [15:11](#)

Another element of any plan is transformative R&D. I do think the administration has focused on that of late. DARPA has a [third wave](#) of AI that they believe is important. I would double down on that class of work to give these systems contextual reasoning and address limitations of today's technology. In the R&D realm, advancing [beyond Moore's law](#) with [quantum](#)

[computing](#) and [neuromorphic computing](#) and other things is also important, as is exploring how AI interacts with institutions, the international order, the labor market. So that's R&D. I'd also include protecting the American worker is critical in any bold AI plan in the United States. You can use things like employer-side subsidies or the earned income tax credit expansion to try to create a more dynamic and resilient workforce.

- BRENDAN: [16:06](#) [Erik Brynjolfsson](#) is a great thinker on this subject. But essentially this allows us to hedge for a wide range of future scenarios while helping people seek out work, lowering the risk of hiring, maintaining the standard of living. There is a need for optimism in the workforce with re-skilling and with lifelong learning and with fair income sharing policies. And then, lastly, our plan should greatly stimulate the development of AI that's ethical and safe. We've continued in that direction in the Department in our AI Strategy, but systems should clearly reflect the U.S. values of freedom and the guarantee of human rights, the rule of law, the stability of institutions, the right to privacy and so forth. We should promote AI standards but also increase the actual market demand for systems that have these properties, that have ethical and safe AI properties. I would love if we took a page from the playbook of clean energy and used the government's purchasing power to try to spur widespread deployment of systems that have these properties. Like fair lending systems or accountable facial recognition or interpretable criminal justice AI. These are all areas where the government can play a huge role.
- CRAIG: [17:56](#) You mentioned a pretty broad spectrum there. Does all of that, in your view, fall under national security or is national security one part of this broader AI strategy that the U.S. should pursue?
- BRENDAN: [18:11](#) I think it's one part. But, I think that the largest gap is actually not in the national security realm.
- CRAIG: [18:17](#) In the economic realm, or the educational real or the political ...
- BRENDAN: [18:24](#) I think it's in economic competitiveness and I think it's in the future of democracy. And I think if these were the two categories of three, the third being more the national security one, I think, with the recent efforts in the Department, with the significant treasure that we allocate to this - though I'm critical of how we allocate it - there is a lot of it there. With the standup of the National Security Commission on AI, you have brilliant minds around the problem, thinking about that and actively at the beginning of a two-year tour to really look at where we should go. So, I think the area for progress, the biggest gap, is

really around economic competitiveness and ensuring the future of democracy.

CRAIG: [19:12](#) Yeah. The gap between where we should be and where we are? Or the gap between what the goals are and what our potential adversaries are doing?

BRENDAN: [19:22](#) I think both. And the way that I would say it is that the potential adversary, for example, China is clearly the example that you brought up in the episode with Trae and Brian, they're using the technology to deepen their grip internally on their populace and to undermine basic human rights in some cases. Additionally, they're using it to spread illiberal practices beyond their borders and to erode public trust in open societies generally, or at least to pose a risk of that. I think that is deeply concerning. But the main focus that I have is actually less around what China can do and is doing. It's more on what the U.S. is not doing and on the vacuum between the two. And I believe that a number of steps that we've taken recently have been only peripherally relevant to that and have not addressed the lack of progress in the United States, have essentially done nothing to close the vacuum.

CRAIG: [20:30](#) Yeah. Just on the Commission, the National Security Commission on AI, how does their mandate fit in, in plugging these holes? Because, are they broad enough to cover the economic competitiveness or the fostering or protecting western style democracy? Or are they really limited to the military piece?

BRENDAN: [20:52](#) I think that group has individuals on it who are certainly thinking about the biggest picture possible. It's chaired by [Eric Schmidt](#) and [Bob Work](#) it's composed of luminaries in relevant fields. And these are folks who are concerned about the biggest questions, around the future of techno democracy versus techno authoritarianism and so forth. I think that they are operating in the solution space of national security topics, and so, their recommendations will naturally hew to that. But they do understand the linkages to the bigger questions. The way that Congress structured that group is they allowed for the pursuit of other topics as necessary. But I know that the National Security Commission on AI will be working with other actors, with the intergovernmental or interagency actors and with the White House to find the right demarcations there.

CRAIG: [21:45](#) Right. when you look at the U.S. compared to China and Russia, who are presumably the top two competitors in this dimension,

how does U.S. stack up against those two countries? Do you have a sense of that?

BRENDAN: [21:58](#) On China, I would comment that earlier this year, I was in several countries in Southeast Asia and I traveled there with U.S. special operations forces. It was a trip that was taken with some senior venture investors and AI experts and it was to advise the head of [SOCOM - Special Operations Command](#) - on steps to take going forward and in light of the changes in emerging technologies, particularly AI.

BRENDAN: [22:23](#) So it was very much going to the forward edge, to the jungles of the Philippines, into areas in remote parts of Thailand and so forth. And what I observed there, being in the region, were a couple of things. One is that China was exceptionally savvy about their public perception and their influence in the region. There were events days prior. There was a bombing at a place called [Jolo church](#) in southern Philippines. And months before that there was a [Thai cave rescue](#). So, we talked to individuals on the U.S. side in the Special Operations Command who had been involved in both of those things. Folks who had come down to help with mission planning or who were actually there in the aftermath, trying to do what they could to save lives in the case of the bombing. And what I would say is that although the U.S. has a sustained presence there and played a significant role in the support effort, and although, in the case of the Philippines, the U.S. contributes over \$100 million, I think, in development assistance and over \$100 million in military aid and is literally [helping fight ISIS-P](#) in the jungle today in southern Philippines, the Chinese involvement there was capitalized on for the purpose of influence in a pretty significant way.

BRENDAN: [23:44](#) And their contribution of about \$700,000 to families was pushed out in social media in a way that showed a level of sophistication around a PR angle. So, in both cases, the Pentagon downplayed the role and I admire that frankly, and that's not the point that I'm trying to make. But public opinion influenced by these operations and by contradictory statements or serious shifts in policy and so forth leaves the people there with difficulty judging whether our nation is good or bad, whether our societal model was good or bad. Another thing that I noticed is, and we talked a lot about, is the strategic character of investment and of industrial penetration of China. One that made the press a couple of days ago was [in Uganda, with Huawei helping to support spying](#). That general exportation of a surveillance stack is very problematic. It's not so much a transfer of capital or a business transaction, but a seizure of power. There, not only is China assuming the major role in some

strategic areas of development in Africa and in other countries, but they are exporting this model at a fairly rapid rate.

BRENDAN: [25:07](#) So, it's a textbook operation that is interesting there. My point is actually that if there are macro strategies available to China around strategic investment, around development,, if this is the playing field, then when you are at the level of an individual military unit, you miss that macro effect almost entirely. You are the parable of the blind men and the elephant and you're staring at a portion of an elephant while the macro strategy is unfolding. So, we need to look and ask ourselves the question of whether we're organized to see these things and to counter them.

MUSIC: [25:52](#)

CRAIG: [26:00](#) In terms of spending, because you know, developing a strategy in the way that you're talking about requires ultimately putting dollars down. What kind of spending is required and is the awareness in the U.S. congress such that they they're likely to put that money down?

BRENDAN: [26:18](#) Well, I would urge just thinking on the order of tens of billions, or a hundred billion dollars over the course of five years. And I say that because, both thinking through this in a bottom up sense and thinking about what are the lines of effort and then what are the orders of magnitude of that and how does it bubble up. And also thinking about reference points around significant mega-projects in the past. And, certainly our way forward would be composed of mega-projects and other spending, non-mega-project spending. To give a reference, the Manhattan project was \$2 billion in yesteryear dollars and that's about 30 billion today, on one well-coordinated project. I think there are lots of different components of what we need to do that that would total around a hundred billion dollars. I will also say, though, that the resourcing inside of the federal government, particularly in the area of defense, is not meager. I mean, it is very significant.

BRENDAN: [27:20](#) It's just distributed badly in some cases. And what I mean by that is that many parts of the Department of Defense are a drag on progress and are in fact doomed by technological change. But they continue to be subsidized and developed. I think maybe 70%, and folks will have to check this figure, of the budget of the Air Force, for example, fits into this category of sustainment. And so, in major paradigm shifts - I think of Europe in the 1960s during the space race. You had countries like France spending their treasure to subsidize their coal industry or

their agriculture, not competing in innovation, not competing for outer space in the way that they could have. And if a fraction of those subsidies to the past were devoted to research or to technology delivery, they could have greatly improved in areas where innovation was crucial. So, I think there is a question around, 'how do you ensure continued economic competitiveness?' I think that's on the order of magnitude of tens of billions, or a hundred billion dollars question. And then I think there's a question of how does the Department look at its resources and reallocate to avoid these subsidies?

CRAIG: [28:44](#) You keep referring to the Department, but is that the right structure to do this? Does the U.S. need a new department, an AI Department, where the money is concentrated and the allocation is coordinated instead of either through the Department of Defense or spread out without much coordination among multiple organizations and initiatives?

BRENDAN: [29:09](#) Well, what I would say is that there's a view that the agencies, the existing agencies, cover down on the problems that we face and that relatively little needs to be done inter-governmentally or at the interstices of those agencies. I don't hold that view. I think that there are problems that we miss almost entirely by focusing only on the agency structure. I think there are reasons why you have these integrative groups like the National Security Council or the Office of Science and Technology Policy. There's another view that says that this is solved principally through a new effort, through a new [Sandia National Lab](#) for AI-type of thing. And I think there is a middle ground where there are incredibly important roles for existing agencies - you know, DARPA focusing on the third wave of AI, the National Science Foundation applying AI to materials science - that are along the lines of effort that I spoke about.

BRENDAN: [30:16](#) But I do think that there is a role for a new structure. I think to ensure sound stewardship you want to have, if for no other reason, a way to bring the best brains to the problem. And so, having an organization that - call it the American AI Institute, for example - that brings together leaders in the public sector with world-leading academics and businesses and philanthropists in a very flexible structure that allows them to coordinate very closely with leading centers around the world. It's the Silicon Valley management truism that [Keith Rabois](#) talks about; you have to get the right people around the problem. And so, in some sense, the National Security Commission on AI is very encouraging in that regard, because, if for nothing else, it's getting brilliant people around the problem. So, I think that makes a lot of sense to have a structure like that, that gives

advice back to government, conducts its own research and does these interfacing roles.

- CRAIG: [31:21](#) But wouldn't it make sense, for example, you used the number \$100 billion over five years, if you want to spend that effectively, shouldn't it be allocated to a central authority and have that authority, that's peopled by the kinds of people you're talking about, coordinate the spending of that money and the prioritization that you talked about?
- BRENDAN: [31:43](#) Yes, I think so. I think it would require a plussing up an existing agency by several hundred and probably even thousands of individuals over time over the course of several years. But I do think for coherence that would be best managed in one structure. I think that it would end up then resulting in other agencies being the means of executing those dollars. But, yes.
- MUSIC: [32:09](#)
- CRAIG: [32:18](#) Can you give, in broad strokes, where you think a hundred billion dollars in five years could be spent effectively?
- BRENDAN: [32:25](#) Well, so let's get specific. I talked about the one goal being increasing the supply of skilled AI professionals. I think if you funded a talent pipeline that tripled the current number of US citizen grad students in computer science and gave them \$25,000 a year, you could add \$1 billion a year to the Department of Education or to the National Science Foundation and cultivate a huge new group of high skilled domestic talent. That's one example. I think you could address the corporate underinvestment in worker training in AI through nearly a billion dollars a year with IRS and Treasury by having a tax deduction for AI worker training within U.S. companies. Other examples are around the access to infrastructure - having an American cloud credit that allowed for essential computational resources to be allocated using a free-market, credit-based approach.
- BRENDAN: [33:21](#) If you had \$10,000 per student per year, that would be \$1 billion and it would essentially have ideas be decoupled from hardware constraints in the economy. Other examples are for \$100 million, on the order of that per year, you could stimulate progress in cyber-physical systems with shared national testbeds. So, imagine if you had a national robotics test bed that can be accessed across the country or if you had test sites to integrate drones into U.S. airspace or if you had an autonomous greenhouse. These are things that big companies can do today. A self-driving car test track on a military base. These are things that large Internet companies can do. But they are a major

barrier to most. And people want the government to be involved with administering those things because it builds competency, creates information for policymakers.

BRENDAN:

[34:24](#)

And we could do that for I think a reasonable amount of money. When you talk about applying AI to global problem solving, I think there is almost no end to the inspirational and daunting tasks that you can apply it to. I mentioned humanitarian assistance and disaster relief. There are so many though that we could formulate. I've seen recently in a collaboration between [AAAS](#) and [Schmidt Futures](#), a catalog of these things, a '[Moonshot Catalog](#).' So, there are a number of interesting problems there. Basic research is expensive. Going beyond Moore's law. You could imagine the quantum computing and neuromorphic computing being well over a billion dollars a year. I will also say that when we talk about using government purchasing power to promote development of AI that's ethical and safe across numerous applications, for \$400 to \$500 million a year, the government purchasing power could start to fill in the market with many different systems that have these beneficial properties in areas where government plays - like safe aviation software and lending systems and facial recognition and criminal justice AI, certainly.

BRENDAN:

[35:30](#)

But even broader than that, I think a huge portion of this would go towards the workforce. You have real uncertainty around the range of future scenarios. The economic benefit could go to those at the top who own the robots. Or there are positive scenarios where wealth is distributed effectively across low, medium and high earners. Or there are other scenarios that are bimodal, or different combinations of those things. The earned income tax credit allows you to nicely deal with those. You need to wind it down or you could spin it up to be able to maintain standards of living across a range of scenarios. So that is probably a billion and a half dollars per year to have some efficacy there and I think would be important, and maybe crucially important, in the coming years. So those are just some examples. I'll throw out one that doesn't have a big number, but is important, which is fostering beneficial alliances and partnerships. There would be a great return on investment. I think strengthening U.S. capacity to set global standards and to establish an international economic and security architecture. I think that's easy to under invest in. And if you have one human being who's focused on that task, it won't get done. But if you staff these emerging technology policy teams and organize them appropriately for small dollars, then they can help bring together, particularly, the democratic countries.

- CRAIG: [37:01](#) What about data because the bottleneck right now is not really so much in the basic research as it is in finding data pools large enough to test and train systems on. Could the U.S. play a role in pooling data or breaking down silos? And in some of the government collected data that's not accessible right now.
- BRENDAN: [37:26](#) Absolutely. And I think in a few different ways. One is unlocking data that exists already in the public sector and across the economy at the federal, state and local government level and in places like healthcare institutions. And so how do you unlock it and use it to create high quality training data, for example? Well, one, you add data engineering resources to government - humans who know how to do that sort of thing, which is lacking. But additionally, you have research and pilots and standards in areas of, broadly speaking, privacy-preserving AI. This happens to be something that I think is really well aligned with our societal model. [Federated machine learning](#) is one example of a technology that would be beneficial here and could be applied in a healthcare context - is being applied, but can be applied broadly.
- BRENDAN: [38:22](#) IP and copyright language is another area. There's a lot of gray area with data on the Internet that could be resolved by the government or by the government working with the private sector. So, one category is around existing data. The other is around new data. Very transformative, and the government has a unique ability to create these datasets and the jobs that go along with them, because frankly, it's a human intensive task most of the time. This data can catalyze new industries. It can spur new breakthroughs. Examples are vehicle crash data or brain imaging data. And we can do this in a way that really broadens accessibility to the data across society. The third piece, though, is, I think, having the United States and universities help in getting the data and making it available, preparing and making available.
- BRENDAN: [39:14](#) So imagine a network of U.S. universities that are operating as data brokers and are helping through government funding to coordinate the identification of data that is relevant to you know, problems, which is constantly changing the data as a function of where the technology is at the frontier. They're helping with the collection, that transformation and sharing of it. [UK biobanks](#) is a good example of this done in a narrower context, in the United Kingdom. But I think universities would love to be enlisted in that effort. And so, you could imagine you have \$300 million a year going to the Department of Education or NSF to establish that brokerage network - and take some of

that load off the government, which frankly isn't equipped to do it well.

CRAIG: [40:02](#) Yeah. We talked briefly before we started the interview about whether addressing climate change with machine learning strategies fits into this. And regardless of what your view is on the divide between what's manmade and what's natural, I don't think anybody argues that there is no global warming. And AI and machine learning can certainly help address that. Do you see that falling within the government mandate or specifically the national security mandate?

CRAIG: [40:33](#) So, I think with climate change there are near term harms that fall into the national security realm. Examples are the impact to refugee flows or the creation of unlivable areas or infectious disease spread or certainly extreme weather events in parts of the world. Arctic passages, I think back to my time on submarines, could open up and that could change the national security calculus in the northern European area. But the major costs of climate change, for national security or otherwise, are still probably several decades away. I think casting it onto the national security landscape, though accurate, is sometimes done to make political coordination more likely or to make it reach those people who might not be persuaded by scientific argument alone. And that's frustrating to me because I think that the meta point is that national security isn't the only vector for progress or for political coordination, for resourcing.

BRENDAN: [40:33](#) Climate change is an issue irrespective of its national security implication. And I believe that there is a chunk of humanity that has the audacity and the determination and ingenuity to go work on this at a massive scale where political coordination has failed. So, this is an area where I would love to see large scale efforts to catalyze progress and even to catalyze the kind of coordination that is needed for the government to do its part. I think we can pull the government in to this one by catalyzing something large. It doesn't solve necessarily for the deficit of political coordination that is present in some global public goods issues outside of climate change. It also doesn't solve for this anti-science sentiment that is present, or that I perceived as present, in government. To some degree, I think we risk having [banners flying and drums beating and marching backward as they say in *Inherit the Wind*](#) about evolution when we're talking about climate change and ignoring the science.

CRAIG: [42:38](#) But if it were put under the national security umbrella, it would get the funding that is needed for this group of people, as you say, that can see that machine learning can have an impact.

Whereas if it's left to a loose coordination among NGOs and academic institutions outside of government, it's never going to get the kind of funding that it needs.

- BRENDAN: [43:02](#) I think in the current political context, that's the case. And I think there are real needs for funding around the issues that I mentioned, the climate change impacts to national security that I mentioned. But the idea of only being able to go at that from a national security line is deeply problematic to me. I think it's a good example of an effort that the United States government, writ large, could play a leadership role in and it could go way beyond national security because the issue goes ...
- CRAIG: [43:31](#) I agree. I just - it's not happening. And, since there is this focus in the last year on coordinating and funding artificial intelligence research and development at a national level, it seems like an opportunity then to get money to that area. It's not just national security, it's critical for humanity.
- MUSIC: [43:55](#)
- CRAIG: [44:02](#) So, a hundred billion dollars over five years, you've given an outline of how that might be spent. Would that kind of funding, properly allocated, keep us competitive with China, given the amount of resources that China is throwing at the problem?
- BRENDAN: [44:18](#) I think the areas that I mentioned would help make immense progress. I think the way that I'll answer this question, though, is separate from the resource level.
- BRENDAN: [44:28](#) So, I think beyond the spending on economic competitiveness and on shoring up the threats to democracy, I think one piece that's missing is one of the reasons that the U.S. was very successful in the period from the Cold War to the first Sputnik. There was this special combination of forces and special fusion of talent across government and private sector and academia that created huge dynamism and progress. We need to find the right units adapted to the current conditions that help us compete here
- CRAIG: [45:01](#) Is Russia a competitor in this field, a credible competitor? Or is it really a lot of bluster.
- CRAIG: [45:10](#) Well, the trend that I'm most concerned with when we talk about Russia is that we're heading towards a very networked world. And you add AI to this networked, digital domain, and everything can be made convincingly synthetic and proliferation of information is the default. So, I think this has huge potential

to disrupt democratic and undemocratic societies alike. I think Russia seems ready to, as a matter of strategy, to seize upon the age of institutionalized fake news and of deep fakes and so forth on social media. So that could be disruptive in a whole range of dimensions, certainly to politics. Also, to financial markets. Big impact there. Niall Ferguson writes well about this with respect to Russia's concern here. The giant networked platforms, the Internet platforms, social media and so forth, have created actually a much less stable public sphere. And so, the fake news and extreme views actually get proliferated in these contexts because of the business models and the algorithms. So, I think Russia is very problematic there and had effort in creating this decisive outcome in the election of the last president. But I think going forward there are many more intended consequences from Russia and unintended consequences as well with this thing.

CRAIG: [46:32](#) You talked about the need to not only increase the human capital domestically, but to increase the flow of qualified researchers into the U.S. from abroad. And there's been this trend that I think is dangerous of a suspicion and resistance toward accepting more, particularly Chinese, researchers and engineers and students in this space because you start pulling that apart and it hardens positions on either side. Whereas, although there are obviously risks in that flow of information back and forth, it tends to bind countries together. How do you feel about that?

BRENDAN: [47:08](#) Well, clearly it can be used for coercive gain, but I think that the entanglement or economic interdependence actually tends to bind us together in a way that promotes stability. I think also Chinese investment activity, and the question is about human capital as well, but the investment activity tends to help, because it increases the efficiency of our businesses and our production. I think going back to Alexander Hamilton, that's been a principle of how we think about foreign investment in the United States. But on that the gift of global talent, Chinese and others, maybe especially Chinese given their position in AI, is clearly ours to lose. It's an immense loss to have these brilliant minds not sticking around after their PhD or going back to China. I always think about the quote of the Russian Grand Designer in *The Right Stuff*, which is, '[our Germans have to be better than their Germans.](#)'

BRENDAN: [48:07](#) And it's kind of a funny phrase, but it gets at this point of, if we can have a country that is the magnet for global talent, this is a very special property. So, I think to your point that, with protection style thinking, it either tends to be irrelevant to the

problem or deals with it peripherally. And I think the problem is not what China can do, but what the U.S. cannot do and what the vacuum is between the two. So, it's either irrelevant to the problem, which is a problem with forward progress, or it's misguided action or it's jingoistic in nature or leads to a bad distinction between the PLA and the Chinese national who's doing a PhD program.

- BRENDAN: [48:52](#) That being said, I will say that there's a lot you could say about Chinese technology transfer and IP theft. There's a lot that's been written about it and uncovered and I do think there's more transparency and information that needs to surface on that. For example, if someone from the PLA applies three or five times under different names to work at a big tech company, I would hope that there are channels to be able to share that, short of calling Sundar and Satya and Mark and Jeff to testify in Congress. That's a real concern. And I think companies sometimes don't feel like they have a channel for sharing that.
- CRAIG: [49:33](#) One of the advantages that China has is the closeness between the government, the military and industry and how there are really no walls there. And how the U.S. government has to fight this public relations battle in order to work openly with industry on certain projects. Project Maven is a perfect example. Do you think it's overstating to say that the U.S. is hobbled by that?
- BRENDAN: [49:59](#) Well, I think people forget that before the Cold War and this first Sputnik period that we talked about, not only was business separate from government, but it was constantly struggling with it. And there was a chasm between intellectuals and businessmen. And this remarkable change happened where we started to fuse talents and create a joint technostructure, like an integrated entity. And I think that the model worked well for several decades. I think we need a new model, a new, call it a biological unit because it's adapted to the current conditions, and we certainly haven't found it. So, I think we are today in a pretty wanting spot as far as that collaboration, the depth of collaboration there.
- BRENDAN: [50:44](#) But I think that's our charge just to find it. With Project Maven, we probably took steps forward and took steps backward on that. But that's the challenge.
- CRAIG: [50:54](#) Yeah. but you don't see it as a particular problem right now?
- BRENDAN: [50:58](#) I do, insofar as the core challenge is to find that new structure to work together. I just am optimistic that we can do it. And I think we can do it based on inspiration and based on being

called forward to solve problems that are important and not based on compelling businesses to do so. I think we can do it through the combination of capital market mechanisms, the government wielding its purchasing power and being a monopsony, the single customer, in some areas, and then also through inspiration based on values. And I think people were ready for that and we haven't done quite a good enough job in inspiring them for some of these important missions.

CRAIG: [51:37](#) You talked about China exporting its surveillance technology to strengthen authoritarian regimes or regimes that are client states, or governments that emulate the Chinese model. How does the U.S. counter that? Trae said something to the effect that the U.S. hasn't stepped up with leadership. But I'm not sure how you step up with leadership through standards and that sort of thing.

BRENDAN: [52:06](#) You can export a different stack, for one. And I think there's really exciting work happening in privacy preserving AI, federated machine learning, that could be proliferated to make for a privacy preserve and value aligned AI solutions and there can be a real big western effort to do so. I also think back to the android model of a free OS that was so good and so useful and so ubiquitous that it reduced the utility of all other non-compliant systems that, for example, in this case, noncompliant could be things that lack off switches or ethics libraries or things like that. So, I think there are interesting head to head stacks or standards that we could promote.

BRENDAN: [53:01](#) And another thing is of course, you can lead in standard setting. And I think the government has started to engage there. But we can engage in and shape the global process a little more, more deeply and then we can also lead in setting some of our own. So, on paper, set standards that are that are world-leading in ethics and safety. But, I basically think that those are the areas we go with. And then I also think there's an interesting counter to the stack of technologies that is surveillance focused and is focused on controlling populations and that is to do something different and to have a much more inspirational set of technologies that you're working with other countries to develop. And I would say around humanitarian assistance and disaster relief is a great example. And what if we had a support and response system for the world in times of disaster that the U.S. helped bring into being and work with academics and companies around the world. I would find that incredibly motivating and it would be a way to mobilize whole governments and whole societies to do some of these things that we want to do, and frankly would be a much better way of

promoting our value system in countering the technical authoritarian stack.

MUSIC: [54:05](#)

CRAIG: [54:13](#)

How much of a threat do you see China as in this regard? I mean, is there a danger that China will get ahead in spreading technology and with it, it's government model or ideology - if there is an ideology anymore - globally, if the U.S. doesn't step in to counter that.

BRENDAN: [54:28](#)

Yes, I absolutely do. There was a French writer who talked about the rise of America and the impact of American industry and technology and culture on Europe in the 1960s. It's called [The American Challenge](#). And he argued that here we are in Europe, we're more prosperous than ever, but we are slipping behind in this push for technological progress and we need to galvanize our forces to confront the challenge. And Europe did not succeed in doing that. They did not work out a coherent strategy that could unite different factions. And they had their outmoded institutions in their entrenched privileges and outdated techniques that hampered them. And they missed the greatest industrial adventure of all time in the space race. America left Europe in the dust. And I think we look at the speed and power and pervasiveness of Chinese investments and activities and AI, and they are a warning to us all around what kind of future do we want. So, I think we have to remaster the art of organization and mobilize people to build this future. But absolutely I think we can do it.

CRAIG: [55:33](#)

Other than Russia and China, who do you see as the next competitor in line?

BRENDAN: [55:39](#)

Well, the National Defense Strategy outlines a couple of other nations and Iran and North Korea clearly comes to mind. But I would also say that with AI in particular, there is a bit of a change in the economics of power or the ability to scale – an 'ensemble of self.' And so, the AI version of an Osama bin Laden where a nation state was going up against an individual in some ways I think is a very real possibility given that they could distribute their impact across the global economy and wreak havoc through threat vectors. So, maybe it's a nation state, but maybe it's others. I know that military thinkers are grappling with that possibility as well.

CRAIG: [56:19](#)

Yeah. This focus that has come about in the last year or so, where do you think the U.S. will be in five years? Do you think

that there will be a robust, coordinated, well-funded national strategy or do you think it's going to take longer than that?

BRENDAN:

[56:36](#)

I think there will be. For all the critique that I shared in this conversation, I think that there is a real chance at bold progress at proceeding in what [David Ignatius](#) called "[the American way](#)," which is not about the national geography, but about a set of values that really came out of the Enlightenment. And I think using our economy, our dominant, vibrant economy, to help resource our plans, I think we can get there. I think there's a huge amount of organizational work needing to be done. And there are people who currently are at the periphery of this who will need to get more involved in the private sector and will need to work on the toughest problems. But I'm optimistic that we can do this. To underscore this in a dramatic way, there were those people who watched the decline of Rome or of Byzantium and caught a glimpse of the future of what was coming. And it wasn't enough to change the course of history. I think the United States can master its fate and can, with other techno democracies, bolster and extend this model. There are risks if we that we don't do that. So, it's a significant call to action for us all.

CRAIG:

[57:53](#)

That's it for this week's podcast. I want to thank Brendan for his time. For those of you who want to go into great depth about the things we talked about today, you can find a transcript of this show in the program notes or at our website, [eye-on.ai](#). The transcript includes relevant links to make it easier to explore. I encourage you to download and read it because I find you discover many things by reading that the ear misses. Let us know if you find the podcast useful and if you have suggestions about how we can improve.

The singularity may not be near, but AI is about to change your world, so pay attention.