JIM BANKS
INDIANA

United States Senate
WASHINGTON, DC 20510

May 14, 2026

The Honorable Scott Bessent
Secretary
Department of the Treasury
1500 Pennsylvania Avenue NW
Washington, DC 20220

The Honorable Sean Cairncross
Director
Office of the National Cyber Director
1600 Pennsylvania Avenue NW
Washington, DC 20500

The Honorable Michael Kratsios
Director
White House Office of Science and Technology Policy
1650 Pennsylvania Avenue NW
Washington, DC 20504

The Honorable Howard Lutnick
Secretary
Department of Commerce
1401 Constitution Avenue NW
Washington, DC 20230

The Honorable Marco Rubio
Secretary
Department of State
2201 C Street NW
Washington, DC 20520

Dear Secretary Bessent, Director Cairncross, Director Kratsios, Secretary Lutnick, and Secretary Rubio:

I am writing to commend you for your efforts in managing the novel security threats posed by advanced artificial intelligence (AI) systems, including new considerations raised by Mythos.

The Trump Administration has been clear about the need for the United States to lead the world in understanding and evaluating the national security implications of advanced AI systems. That need has only become more urgent in light of Mythos, an AI system with advanced cyber capabilities. This is an important moment for the country: Mythos underscores that AI will have profound implications for cybersecurity and national security.

As President Trump and administration officials engage with counterparts from the People's Republic of China (PRC) this week, it is wholly appropriate for AI to be on the agenda. While some have naively argued the U.S. should pursue an AI "grand bargain" or arms control analogue with China, **I encourage the administration to consider a less ambitious, narrower range of AI issues to engage with the Chinese.**

Importantly, Mythos is only one example of a much broader challenge. **While cyber capabilities are the focus today, the conversation must extend well beyond cyber.** AI is improving

rapidly. Advanced AI systems are expected to develop increasingly consequential capabilities across military, intelligence, biosecurity, and other national security domains. As you know, this is why it is imperative that the United States maintain and extend its lead over the People's Republic of China (PRC) in advanced AI development.

During the Cold War, atomic weapons introduced a paradigm shift in the geopolitical world order. AI has the potential to do the same. **There may be certain critical thresholds in AI development that are imperative to reach before adversaries.** One example—potentially the most meaningful threshold articulated thus far—would be AI systems that *outperform humans in making new breakthroughs in AI and developing increasingly more powerful AI systems*. AI systems with such research and development capabilities would potentially lead to rapid and self-reinforcing improvements: America must reach this threshold before the PRC.

**We must also recognize that advanced AI presents serious risks.** In the wrong hands, AI can be weaponized against the American people. More broadly, there are unresolved questions about *whether we currently have the ability to keep powerful AI systems under our control*. These questions are especially salient as companies pursue artificial superintelligence—AI models that can do anything a human can do but much, much better. Our major AI companies believe superintelligence could pose significant risks that our current scientific understanding is not yet capable of addressing.

**America must embrace the challenging dual mandate of winning the AI race against the PRC while navigating critical security challenges along the way.** This is what makes AI policy a particularly difficult domain. Some have called for a laissez faire approach to AI, which neglects to consider the serious and novel threats posed by technology. Others call for extremely heavy-handed approaches that would cede our lead to the PRC. Still others have called for pursuing bilateral AI arms control agreements with the PRC.

Such calls for bilateral arms control-style agreements ignore the CCP's long history of deception and non-compliance with its international commitments. They also ignore difficulties in monitoring and compliance compared to nuclear arms control agreements—which our adversaries have frequently violated anyways.

**However, there may be limited areas where it does make sense to engage in dialogue with Chinese officials.** I would recommend a simple test: if it is in America's national interest to adopt a given AI policy unilaterally due to security concerns, *even if we knew the PRC may cheat*, then it is worth engaging the Chinese on the possibility of reciprocal action.

I applaud the Administration for deliberating carefully about AI policy and ensuring that we do not fall into either extreme. Mythos provides an excellent opportunity for the Administration to revisit key questions about AI policy, semiconductor policy, the relationship between the government and the private sector, and the broader U.S-China competition.

With this in mind, I recommend considering the following questions as you continue to focus on AI and national security topics:

1. **Infrastructure and expertise.** What infrastructure and expertise are needed to ensure the United States leads in evaluating the national security implications of advanced AI systems? How do we ensure that we have the insight required to assess models not only for cyber capabilities but also in areas like military applications, loss of control to AI systems themselves, automated AI research and development capabilities, and other security domains?

2. **Information-sharing requirements.** Under what circumstances should frontier AI companies be expected or mandated to report national security information to the government? Under what circumstances should they be expected or mandated to provide models, the results of model evaluations, or other critical information? What information might be essential for the government to know even before models are publicly or externally deployed?

3. **Incident reporting**. What kinds of incidents should frontier AI companies be expected or mandated to report to the U.S. government immediately upon discovering them? For example, how should companies notify the government if the PRC steals the weights of an advanced AI system, if an AI system attempts to exfiltrate its own model weights, or if models engage in unexpected behavior that poses risks to U.S. national security?

4. **Limitations of voluntary agreements**. How robust are voluntary agreements around information-sharing and model access? Suppose a company—for financial or political reasons—decided not to share relevant national security information or model access with the United States governments. What kinds of policies are needed to address these gaps?

5. **Policies for PRC engagement**. What is the universe of AI policies that we would want to pursue unilaterally due to security concerns, that we should engage with the PRC on—even if we knew they would cheat?

6. **Support from Congress.** How can Congress support the Administration as it continues to explore new policy approaches relating to AI and national security?

Mythos presents us with an opportunity to think about the future of AI, its national security implications, the stakes of the AI race with the PRC, and the appropriate relationship between the government and the private sector. To support your work, I request a staff-level briefing within 60 days to discuss these topics.

Sincerely,

Jim Banks
U.S. Senator for Indiana