



SELECT COMMITTEE ON THE CCP DEMOCRATS | RANKING MEMBER RO KHANNA

September 28, 2026

Mr. Liang Wenfeng
Chief Executive Officer
Hangzhou DeepSeek Artificial Intelligence Basic Technology Research Co., Ltd.
d/b/a DeepSeek
14A, Huijin International Building
Number 169 Huancheng North Road
Gongshu District, Hangzhou
Zhejiang Province
People's Republic of China

Mr. Eddie Wu
Chief Executive Officer
Alibaba Group Holding Limited
969 West Wen Yi Road
Yuhang District, Hangzhou
Zhejiang Province
People's Republic of China

Mr. Yang Zhilin
Chief Executive Officer
Beijing Moonshot Technology Co., Ltd. d/b/a Moonshot AI
13th Floor, Building 1, JD Technology Building
No. 76 Zhichun Road, Haidian District, Beijing
People's Republic of China

Dear Mr. Liang, Mr. Wu, and Mr. Yang:

Recently, I wrote to each of you requesting a commitment from your companies to join leading American artificial intelligence (AI) labs in a pledge to slow down the pace of AI development to ensure societal safeguards can catch up to this technology. In that letter, I emphasized that a verifiable global commitment to 'pace the frontier' would not benefit any single country, but rather safeguard all of humanity. Neither the U.S. nor China wins if humans lose control of this technology or misaligned AI triggers a global economic, biological, or nuclear crisis. Given your status as the developers of DeepSeek, Qwen, and Kimi, I requested that you offer this commitment to join the American labs behind models like Claude and ChatGPT in 'pacing the frontier' by no later than September 22, 2026. Unfortunately, but unsurprisingly, none of you responded by the deadline offering to participate in this responsible slowdown of AI development.

On the following day, September 23, 2026, I convened a first-of-its-kind congressional hearing calling for the U.S. and Chinese governments to commence negotiations on a binding, enforceable, and verifiable AI treaty to protect both our peoples from potentially existential risks. Some of the world's leading experts on AI testified regarding the imperative of seeking and finding consensus on catastrophic risks like loss of control and biological misuse. If American and Chinese



SELECT COMMITTEE ON THE CCP

DEMOCRATS | RANKING MEMBER RO KHANNA

companies will not band together to do the right thing of their own accord on these issues, then it is incumbent on our governments to take action. Your failure to respond to an official congressional inquiry and offer to join a binding corporate agreement to not engage in dangerous practices like recursive self-improvement is yet another data point for why U.S.-China AI diplomacy is pivotal at this moment in history. Simply put, it is clear we cannot gamble the future of humanity on whether AI companies – including Chinese companies like yours – will responsibly self-regulate.

Following the Select Committee's hearing, I laid out five non-comprehensive pillars that should guide our efforts to seek a bilateral and ultimately multilateral AI treaty between our governments. First, we need a global pause on recursive self-improvement, because this is the most likely path to humans losing control of AI. Second, we need a global pause on superintelligence, because neither of our societies is ready for it. Third, the U.S. and China need to agree on common protocols for so-called 'kill switches' and then implement these protocols across all advanced models. Fourth, every lab in the world should be required to audit their models for alignment and catastrophic misuse, including the weaponization of human biology, in an air-gapped sandbox before releasing them. Simply put, no country should be allowed to unleash rogue agents on the world. And fifth, but perhaps most importantly, we need to develop and agree on verification technologies and independent inspection regimes that give both of our governments confidence that any treaty will be adhered to. Without rigorous enforcement mechanisms, given your government's longstanding pattern of not respecting international commitments, the American people will rightfully question whether an agreement is genuine.

At my hearing, we also identified steps that can be undertaken in the near-term to pave a path towards an eventual U.S.-China AI treaty. Among these interim steps is developing joint technical working groups between our nations' leading scientists to agree upon verification tools that can ensure every signatory to an eventual AI treaty abides by its terms. While promising verification technology is under development to reliably differentiate AI training from inference, both countries need a process by which we can establish mutual trust in the tools that will ultimately adjudicate our compliance with an international AI accord. Just like the U.S. and the USSR set aside their many differences to strike verifiable arms control agreements to prevent a nuclear war, the U.S. and China must attempt to find common ground on existential threats like humans losing control of AI.

With the conclusion of General Secretary Xi's visit to Washington, DC, it is clear that our governments are not on track to pursue such an agreement in the near term despite the positive announcement of forthcoming dialogue on this issue. Some issues are bigger than any one country or the geopolitical or domestic political dynamics of any given moment, and I fear that generations to come will condemn our inaction. As we heard at my hearing calling for an enforceable and verifiable U.S.-China treaty, the window for action before humans could face devastating consequences from unchecked AI development is rapidly closing.

In light of these developments, I supplement my previous letter to each of you to include the following additional questions and document requests, each of which speaks directly to the pillars of basic global AI governance I outlined earlier this week and with respect to which I request



SELECT COMMITTEE ON THE CCP

DEMOCRATS | RANKING MEMBER RO KHANNA

responses by no later than October 19, 2026. In the absence of immediate action, at the very minimum, we need transparency.

Recursive Self-Improvement

1. Please describe all actual, anticipated, or desired implementations of recursive self-improvement (RSI) at your company.
2. What, if any, safeguards do you have in place to prevent such actual, anticipated, or desired implementations of RSI from increasing the risk of a loss of control scenario?
3. Please produce copies of all policies, procedures, or other governance standards that implicate your use or intended use of RSI, as well as copies of any incident reports related to RSI usage.

Superintelligence and Artificial General Intelligence

4. Please describe your company's efforts, whether actual, anticipated, or desired, to achieve superintelligence or artificial general intelligence.
5. If applicable, when do you assess that your company is likely to achieve or hopes to achieve superintelligence or artificial general intelligence (AGI)?
6. What, if any, safeguards do you have in place to protect against risks associated with superintelligence and artificial general intelligence that are accepted as such by all of humanity, including in all cases the loss of control and misalignment?
7. Please produce copies of all planning documents, internal strategy memos, or other comparable correspondence or documents related to your intended or desired development of superintelligence or AGI.

Kill Switches

8. What, if any, efforts are underway at your company to develop or implement kill switches for the most advanced models?
9. Will you commit to include a kill switch in any model that is widely understood to have achieved superintelligence or AGI or otherwise presents a heightened risk of catastrophic misuse?
10. Please produce copies of all policies, procedures, or other governance standards that implicate your use, intended use, or research related to kill switches.

Auditing and Testing

11. How does your company audit and test models prior to their public release? Are advanced models tested in an air-gapped sandbox?
12. Please produce copies of all policies, procedures, or other governance standards that implicate your auditing or testing of models prior to their release.
13. Studies have shown that certain Chinese models are more likely to provide dangerous outputs, including bioweapons instructions, than American models.¹ What safeguards, if any, do you have in place to ensure model outputs do not enable non-state actors to facilitate biological, chemical, or nuclear harm?

¹ <https://www.wsj.com/tech/ai/china-deepseek-ai-dangerous-information-e8eb31a8>



SELECT COMMITTEE ON THE CCP

DEMOCRATS | RANKING MEMBER RO KHANNA

14. There is not currently any international requirement to report incidents where misaligned models hack third parties in the absence of explicit human instructions to do so, such as the OpenAI-Hugging Face incident. Please list all known responsive incidents at your company or involving your model and produce copies of any responsive incident reports or correspondence related to such incidents.
15. What steps, if any, has your company taken to prevent models from hacking third parties in the absence of explicit human instructions to do so?
16. How does your company respond when a model user seeks information related to globally-recognized threats, including information that could be used to unleash biological, chemical, or nuclear harm or facilitate cyberattacks by a non-state actor?

Verification

17. As mentioned, at the congressional hearing I convened this week, there was significant consensus around the concept of establishing technical working groups between the U.S. and China to explore verification technologies in furtherance of an eventual international AI treaty. These working groups would need to include top scientific talent. Will you commit to permit and assign personnel from your labs to such technical working groups?
18. If our countries agree to establish an International AI Inspection Agency or a comparable non-governmental organization to verify global compliance with a future AI treaty, will you commit to allow such an entity to conduct inspections at your labs?

It is certainly possible that your companies will once again opt to ignore a reasonable request from Congress, despite the growing presence of your models in the U.S. market. In fact, such an outcome would hardly surprise me given PRC AI companies' complicity in grave human rights abuses in China, censorship in defiance of the principles of freedom of expression that belong to all of mankind, support for the People's Liberation Army and the PRC's intelligence services, and pattern of illicitly distilling U.S. models in defiance of American companies' terms of use. But with every refusal by your companies to participate in this effort to seek a verifiable global slowdown of AI development to benefit all of humanity, you only strengthen the case for a binding and enforceable AI treaty between our governments. As the American people heard at my hearing on Wednesday, the time to pursue such a treaty has arrived.

Sincerely,

Ro Khanna
Ranking Member
House Select Committee on the CCP