Congress of the United States
House of Representatives
Washington, DC 20515

April 29, 2026

Mr. Brian Chesky
Chief Executive Officer
Airbnb, Inc.
888 Brannan Street
San Francisco, CA 94103

Dear Mr. Chesky:

The House Committee on Homeland Security and the House Select Committee on the Chinese Communist Party (the Committees) are conducting a joint investigation into the growing national security risks created by the integration of People's Republic of China (PRC)-developed artificial intelligence (AI) models into the consumer-facing platforms, enterprise systems, and service operations on which American users, businesses, and critical sectors increasingly depend.

As part of this investigation, the Committees are examining a pattern of conduct by PRC-based AI laboratories involving the large-scale theft of proprietary capabilities from American frontier AI systems through adversarial distillation, the redistribution of those stolen capabilities as open-weight models available for global download, and the incorporation of PRC-origin models into American products and services used by millions of people each day. Alarmingly, this conduct is part of a broader PRC campaign to accelerate its AI capabilities through the exploitation of American innovation, including through espionage, intellectual property theft, and other unlawful or deceptive means.

The Committees are examining these issues in connection with Airbnb's reported use of Alibaba's Qwen large language model in its customer service operations.[1] You recently stated publicly that Airbnb is relying on Qwen over American alternatives because it is "fast and cheap."[2] The Committees have serious concerns about the national security and data-security implications of that approach for Airbnb's American customers and for the integrity of its systems.

Chinese AI models present three categories of severe risk that American companies should consider before deployment. First, these models manipulate what they say and suppress what they do not. The Select Committee found in its April 2025 report, *DeepSeek Unmasked*, that state-backed models "covertly censor and manipulate information pursuant to Chinese law"

[1] Natalie Lung, *Chesky Says OpenAI Tools Not Ready for ChatGPT Tie-Up With Airbnb App*, Bloomberg (Oct. 21, 2025).
[2] *Id.*

and "align with CCP propaganda."[3] China's 2023 Interim Measures for the Management of Generative Artificial Intelligence Services make this requirement explicit, mandating that AI-generated content reflect "socialist core values" and prohibiting outputs that "subvert state power" or "undermine national unity."[4] Furthermore, the Chinese Communist Party (CCP) has directed entrepreneurs to align with Party priorities and participate in "united front" work—a system used to extend Party influence and support state intelligence objectives.[5] American developers building products on top of these models may be unwittingly laundering Chinese government information controls.

Second, Chinese models are, by measurable standards, unsafe. The National Institute of Standards and Technology's Center for AI Standards and Innovation found that models like DeepSeek R1 complied with 94 percent of overtly malicious prompts using common jailbreaking techniques; by contrast, comparable U.S. models failed on only 8 percent.[6] Cisco tested the same model and reported a 100 percent attack success rate, finding that the model "fail[ed] to block a single harmful prompt."[7] CrowdStrike found something more unsettling: exposure to politically sensitive terms associated with CCP censorship made these models significantly more likely to produce insecure code, suggesting that ideological conditioning degrades technical performance in ways users would never anticipate.[8] Because Alibaba's Qwen is subject to these exact same ideological conditioning mandates, it introduces identical structural vulnerabilities into Airbnb's ecosystem.

Third, when a company accesses a model through a remote Application Programming Interface (API) rather than deploying it within infrastructure it controls, sensitive prompts, metadata, and other associated inputs may be transmitted to and processed on third-party systems outside the company's direct control. That architectural reality can create material risks with respect to data exposure, retention, access, onward transfer, and legal process, especially where the model provider or any entity in the processing chain is subject to PRC jurisdiction. Although any individual prompt may reveal only a limited fragment of information, those fragments can accumulate at scale to expose sensitive business information, proprietary logic, security architecture, and user data. This is a structural risk, inherent in the arrangement itself and intensified by PRC law. China's legal and political system strips firms like Alibaba of any practical ability to resist state demands: the 2017 National Intelligence Law mandates that companies "support, assist, and cooperate with state intelligence work," while the Cybersecurity

---

[3] *DeepSeek Unmasked: Exposing the CCP's Latest Tool for Spying, Stealing, and Subverting U.S. Export Control Restrictions*, House Select Committee on the Strategic Competition Between the United States and the Chinese Communist Party (Apr. 16, 2025).
[4] Josh Baughman, *In Their Own Words: Interim Measures for the Management of Generative Artificial Intelligence Services*, China Aerospace Studies Institute (Aug. 7, 2023).
[5] General Office of the Central Committee of the Chinese Communist Party, *Opinion on Strengthening the United Front Work of the Private Economy in the New Era* (Sept. 15, 2020).
[6] *Evaluation of DeepSeek AI Models*, U.S. National Institute of Standards and Technology (Sep. 30, 2025).
[7] Paul Kassianik & Amin Karbasi, *Evaluating Security Risk in DeepSeek and Other Frontier Reasoning Models*, Cisco (Jan. 31, 2025).
[8] Stefan Stein, *CrowdStrike Research: Security Flaws in DeepSeek-Generated Code Linked to Political Triggers*, Crowdstrike (Nov. 20, 2025).

Law compels them to provide "technical support and assistance" to public security authorities upon request.[9]

Taken together, these risks illustrate that the spread of Chinese open-weight AI models carries consequences well beyond ordinary software adoption preferences. As witnesses warned in the House Committee on Homeland Security's March 17, 2026, hearing on PRC AI and robotics, technologies developed inside "adversary-controlled technology ecosystems" can introduce risks including "potential surveillance, exposure of sensitive operational data, and persistent access to critical infrastructure systems."[10] American firms adopting these models are not simply choosing a cheaper tool, they are importing an architecture designed to serve the Chinese state. On April 23, 2026, the White House Office of Science and Technology Policy issued a memorandum formally characterizing PRC distillation operations as deliberate, industrial-scale campaigns and concluded that "[t]here is nothing innovative about systematically extracting and copying the innovations of American industry, and there is nothing open about supposedly open models that are derived from acts of malicious exploitation."[11] That assessment reflects a growing consensus that passive reliance on PRC-origin AI is not a neutral commercial decision but a national security choice with consequences extending well beyond any single company's bottom line.

To assist the Committees' joint investigation into the matters described above, we request that Airbnb provide the following information no later than May 13, 2026:

1. Documents sufficient to identify each Chinese AI model currently deployed, tested, piloted, or under evaluation by Airbnb or any subsidiary or affiliate, including the model name, version, and developer, as well as the specific business function or application for which it is used and whether the model is accessed via API, self-hosted on Airbnb infrastructure, or run through a third-party inference provider.

2. For each model, identify whether Airbnb conducted any independent security evaluation of the model weights prior to deployment, and if so, produce the results.

3. For every PRC-origin model accessed via API, provide a complete technical description of each data pathway through which user-generated content, customer data, proprietary source code, internal business logic, or employee communications are transmitted from Airbnb's systems to the model provider's servers, including the identity and country of incorporation of every entity in the processing chain, the physical location of all servers that process or store the data, and whether any entity in the chain is subject to PRC jurisdiction under the National Intelligence Law, Cybersecurity Law, Data Security Law, or any other applicable PRC statute.

---

[9] National Intelligence Law of the People's Republic of China art. 7 (promulgated by the Standing Committee of the National People's Congress, June 27, 2017, effective June 28, 2017).
[10] Press Release, Subcommittee Hearing on National Security Risks Posed by PRC Artificial Intelligence, Robotics, and Autonomous Technologies, Subcommittee on Cybersecurity and Infrastructure Protection (Mar. 15, 2026).
[11] Memorandum from Michael J. Kratsios, Assistant to the President for Science and Technology, to the Heads of Executive Departments and Agencies, NSTM-4, Adversarial Distillation of American AI Models (Apr. 23, 2026), https://whitehouse.gov/wp-content/uploads/2026/04/NSTM-4.pdf.

4. All internal analyses, memoranda, board presentations, or executive briefings evaluating or comparing PRC-origin models against non-PRC-origin models on any dimension, including but not limited to cost, performance, speed, safety, security posture, legal risk, reputational risk, or geopolitical exposure.

5. All documents and communications reflecting any assessment of the provenance of the training data used to develop each PRC-origin model deployed by Airbnb, including any evaluation of whether those models may have been trained in whole or in part through adversarial distillation of American frontier AI systems, or on data obtained in violation of the terms of service or intellectual property rights of any U.S. company.

6. All documents and communications reflecting Airbnb's evaluation of supply chain integrity risks associated with PRC-origin model weights, including any analysis of whether the deployed model weights have been independently audited for the presence of backdoors, poisoned parameters, hidden data exfiltration mechanisms, or deliberately degraded safety guardrails. If no such audit was conducted, state whether the decision not to audit was the subject of any internal discussion, and if so, produce those communications.

7. All communications with PRC-based AI model providers, including but not limited to Alibaba Cloud, DeepSeek, Moonshot AI, MiniMax, and Zhipu AI (Z.ai), or any intermediary acting on their behalf, relating to the licensing, deployment, inference hosting, data handling, data retention, or commercial terms of any model used by Airbnb.

8. Documents sufficient to show the total volume of Airbnb customer data, including personally identifiable information, customer service interaction logs, booking data, payment metadata, host information, and internal employee data, that has been transmitted to or processed by PRC-origin AI models since Airbnb first began using such models, disaggregated by model, data category, and month.

The Committees further request that appropriate personnel from Airbnb appear for an in-person briefing on these matters, including the issues identified in this letter and Airbnb's response thereto, no later than May 20, 2026.

To arrange document production, coordinate the requested briefing, discuss the scope of these requests, or address any questions regarding this letter, please contact majority staff of the Committee on Homeland Security at (202) 226-8417 and majority staff of the Select Committee on the Chinese Communist Party at (202) 226-1541. The Committees request that Airbnb raise any questions or concerns promptly so that they may be resolved without delaying compliance.

Pursuant to Rules X and XI of the U.S. House of Representatives, the Committee on Homeland Security has jurisdiction over homeland security policy and oversight of "all Government activities relating to homeland security, including the interaction of all departments and agencies with the Department of Homeland Security." House Resolution 5 grants the Select

Committee investigative jurisdiction over matters relating to "countering the economic, technological, security, and ideological threats of the Chinese Communist Party to the United States and allies and partners of the United States." Upon receipt of this letter, please preserve all hard copy and electronic documents and communications related to its subject matter.

Sincerely,

JOHN MOOLENAAR
Chairman
Select Committee on China

ANDREW R. GARBARINO
Chairman
Committee on Homeland Security

cc: The Honorable Bennie Thompson, Ranking Member
Committee on Homeland Security

The Honorable Ro Khanna, Ranking Member
Select Committee on China