United States Court of Appeals

FOR THE DISTRICT OF COLUMBIA CIRCUIT

Argued May 19, 2026 Decided September 25, 2026

No. 26-1049

ANTHROPIC PBC,
PETITIONER

v.

UNITED STATES DEPARTMENT OF WAR AND PETER B. HEGSETH, IN HIS OFFICIAL CAPACITY AS SECRETARY OF WAR,
RESPONDENTS

Consolidated with 26-1162

On Petitions for Review of an Agency Action of the Department of War

---

Kelly P. Dunbar argued the cause for petitioner. With him on the briefs were *Joshua A. Geltzer*, *Kevin M. Lamb*, *Anneke Dunbar-Gronke*, and *Megan O. Gardner*.

Tim Hwang was on the brief for *amici curiae* Foundation for American Innovation, et al. in support of petitioner.

Barbara Smith Tyson was on the brief for *amicus curiae* Taxpayers Protection Alliance Foundation in support of petitioner.

Benjamin Klubes was on the brief for *amici curiae* Catholic Moral Theologians and Ethicists in support of petitioner.

Harold Hongju Koh, *Bruce Swartz*, *Alexis Loeb*, *Anthony Schoenberg*, and *John Ugai* were on the brief for *amici curiae* Former Senior National Security Government Officials in support of petitioner.

Ori Lev was on the brief for *amici curiae* Employees of OpenAI and Google in their personal capacities in support of petitioner.

Samir Jain, *Ashley Gorski*, and *Patrick Toomey* were on the brief for *amici curiae* American Civil Liberties Union and Center for Democracy and Technology in support of petitioner.

Brian Scarpelli was on the brief for *amicus curiae* the Association for Competitive Technology (ACT) in support of petitioner.

Matthew Klapper, *Elizabeth Deutsch*, and *Andrew Cherry* were on the brief for *amici curiae* Former Secretary of Defense Leon Panetta and the Institute for Security and Technology in support of petitioner.

Norman L. Eisen, *Stephen A. Jonas*, *Gregg J. Costa*, *Sophia Brill*, and *Connor P. Mui* were on the brief for *amici curiae* 149 Former Judges and Democracy Defenders Fund in support of petitioner.

Sarah Grant, *Sopen B. Shah*, and *Addison W. Bennett* were on the brief for *amici curiae* the Foundation for Individual Rights and Expression, et al. in support of petitioner.

Rakesh Kilaru was on the brief for *amicus curiae* Faith Family Technology Network in support of petitioner.

Daniel W. Wolff was on the brief for *amici curiae* Industry Trade Associations in support of petitioner.

Elisabeth S. Theodore, *Benjamin C. Mizer*, *Samuel F. Callahan*, and *Aaron X. Sobel* were on the brief for *amicus curiae* Professor Alan Z. Rozenshtein in support of petitioner.

Sarah E. Harrington, *Alexander A. Berengaut*, *David M. Zionts*, *Megan A. Crowley*, and *Mishi Jain* were on the brief for *amici curiae* Former Service Secretaries and Retired Senior Military Officers in support of petitioner.

Josephine K. Petrick and *Hayley Landman* were on the brief for *amici curiae* Freedom Economy Business Association and Values-Led Investors in support of petitioner. *George C. Harris* entered an appearance.

Sharon Swingle, Attorney, U.S. Department of Justice, argued the cause for respondents. With her on the briefs were *Brett A. Shumate*, Assistant Attorney General, *Eric D. McArthur*, Deputy Assistant Attorney General, and *Sean R. Janda* and *Brian J. Springer*, Attorneys.

Gina D'Andrea, *Andrew Zimmitti*, and *Joel Thayer* were on the brief for *amicus curiae* Joel Thayer, Senior Fellow at the America First Policy Institute in support of respondents.

Before: HENDERSON, KATSAS, and RAO, *Circuit Judges*.

Opinion for the Court filed by *Circuit Judge* KATSAS.

Dissenting opinion filed by *Circuit Judge* HENDERSON.

KATSAS, *Circuit Judge*: This case arises from a decision by the Department of War to exclude Claude, an artificial-intelligence product developed by petitioner Anthropic PBC, from its supply chain under the Federal Acquisition Supply Chain Security Act of 2018. The Department made this decision after Anthropic refused to relax contractual prohibitions on the use of Claude for lethal autonomous warfare or domestic surveillance. Anthropic challenges the exclusion as arbitrary, unauthorized by the governing statute, and unconstitutional.

We reject these challenges. The Department had ample support for its conclusion that the continued integration of Claude into the Department's information systems, by the Department or its contractors, presented a statutorily covered national-security risk. As Anthropic admits, the company encodes restrictions into Claude that prevent the model from performing tasks that Anthropic wishes to prevent. On more than one occasion, these restrictions have stopped Claude from performing tasks requested by government users. And recently, a dispute arose over whether the contractual prohibitions barred the use of Claude in an ongoing overseas military operation, leaving the Department uncertain whether Claude would perform as needed and intended.

Anthropic's constitutional claims are also without merit. Its due-process claim fails because the Department promptly notified the company of the exclusion and its supporting rationale, and then gave the company a fair opportunity to contest the exclusion. And Anthropic's First Amendment

claim fails because the Department excluded Anthropic from its supply chain based on the company's refusal to assent to a contract term that the Department deemed essential, not based on the company's support for greater governmental regulation of AI technology.

I

The Federal Acquisition Supply Chain Security Act of 2018 authorizes "covered procurement action[s]" to prevent agencies from using information technologies that pose a risk to national security. 41 U.S.C. § 4713(a). Such procurement actions include barring agency contracts with a particular supplier and subcontracts that use the supplier to perform work for the agency. *Id.* § 4713(k)(4)(A), (B), (D).

To take a covered procurement action, an agency head must first make a written determination that use of the authority to take the action "is necessary to protect national security by reducing supply chain risk" and that "less intrusive measures are not reasonably available" to reduce that risk. 41 U.S.C. § 4713(b)(3)(A), (B). The determination must also specify the class of covered procurement actions the agency may take under it. *Id.* § 4713(b)(3)(C). The statute defines "supply chain risk" to mean "the risk that any person may sabotage, maliciously introduce unwanted function, extract data, or otherwise manipulate the design, integrity, manufacturing, production, distribution, installation, operation, maintenance, disposition, or retirement" of covered information-technology products "so as to surveil, deny, disrupt, or otherwise manipulate the function, use, or operation of" those products or the information stored or transmitted on them. *Id.* § 4713(k)(6).

Before determining that it is necessary to take covered procurement actions, the agency head must provide the

supplier with notice and an opportunity to respond. 41 U.S.C. § 4713(b)(2). However, the agency head may "temporarily delay" providing notice if he "determines that an urgent national security interest requires the immediate exercise of the authority." *Id.* § 4713(c), (c)(1)(A). In that instance, the agency head must provide notice and an opportunity to respond "as soon as practicable after addressing the urgent national security interest." *Id.* § 4713(c)(2), (2)(A). Once the agency head makes a determination, the agency may take all procurement actions covered by the determination. *Id.* § 4713(a), (b)(3)(C).

The statute channels judicial review of covered procurement actions into this Court. It provides that a party notified of a "covered procurement action under section 4713" may file a petition for review of that action in this Court within 60 days of the notification. 41 U.S.C. § 1327(b)(1). The statute also bars other judicial review of any "action taken under" section 4713. *Id*. § 1327(a).

II

A

Anthropic develops Claude, a family of artificial-intelligence models. Claude utilizes large-language models, which are algorithms "trained on massive datasets to identify patterns and associations in language." App. 6. Claude can "respond[] to a wide range of user inputs, or 'prompts,' in an intelligent, human-like manner." *Id.* It "can even act autonomously, executing tasks without requiring ongoing user direction." *Id.* at 7. In the context of warfare, this means that Claude could be deployed in a way that "independently identifies and classifies an object as a military target, determines engagement criteria are satisfied, and launches a

weapon strike." *Id.* at 16. The power of this new technology is obvious—as is its potential for misuse.

To reduce the risk of misuse, Anthropic employs three distinct kinds of restrictions. First, it "seek[s] to embed safety considerations directly into the model itself." App. 8. For example, Anthropic has disabled Claude from performing specific tasks such as making biological, chemical, nuclear, or radiological weapons. *Id.* at 103. Beyond that, Anthropic trains Claude to conform to a constitution developed for it by Anthropic. According to Jared Kaplan, Anthropic's co-founder and Chief Science Officer, this constitutional training makes Claude follow "a set of normative principles, like balancing helpfulness against harm avoidance, and respecting values such as individual privacy and political freedom." *Id.* at 8. Dario Amodei, Anthropic's Chief Executive Officer, explains that this training, focused on "high-level principles and values," imbues Claude with an "identity, character, values, and personality" that lead to what Anthropic deems "a coherent, wholesome, and balanced psychology." *Id.* at 93–94. Anthropic considers this kind of "model development" to be "at the core of [its] mission." *Id.* at 2.

Second, Anthropic builds into its products "technical measures that stack on top of the model itself." App. 8. These measures include monitoring systems to detect harmful activity and targeted interventions to prevent it. *Id.* Anthropic began to develop this "second line of defense" around mid-2025, because "all models can be jailbroken." *Id.* at 103.

Third, Anthropic contractually prohibits uses of Claude that it deems inappropriate. As summarized by Kaplan, its Usage Policy prohibits "unacceptable" uses including "surveillance, compromising computer systems or networks, and designing weapons." App. 9. The actual Usage Policy

imposed by Anthropic on Palantir Technologies, Inc., which analyzes data for the Department of War, is more extensive. Among other things, it prohibits Palantir from using Claude to compromise children's safety, incite violence or hateful behavior, invade privacy, create emotionally harmful content, spread misinformation, interfere in elections, or monitor individuals' physical locations. *Id.* at 394–400. Anthropic informs us that such usage restrictions reflect "the very purpose for which our company was founded," and their removal would "contradict our deeply held values." *Id.* at 10.

B

Over the past two years, the Department of Defense (which now calls itself the Department of War) has greatly expanded its use of AI. During this time, Anthropic partially—but not completely—relaxed its use restrictions to accommodate the Department.

In 2024, the Department and the intelligence community began using standard, commercially available models of Claude in their classified systems, working through contractors with access to Claude. App. 279. Consistent with Anthropic's model training, Claude "refuse[d]" to perform "tasks that were appropriate in a national security context—such as summarizing threat assessments, processing classified documents, or translating intercepted materials describing violence." *Id.* at 280.

In response to this problem, Anthropic developed a special "Claude Gov" model that would perform such tasks for the national-security agencies, which it released in March 2025. App. 280–81. Anthropic also developed a "government-specific addendum" to its Usage Policy, which contractually allowed certain uses that it would deny to private customers. *Id.* at 12. Anthropic describes this addendum as "designed to

strike a balance between enabling national security beneficial uses and mitigating potential harms." *Id.* Over time, Anthropic came to permit the Department to use Claude to design more effective weapon systems, to analyze foreign intelligence, and to conduct offensive cyber operations. *Id.* at 12, 278. However, Anthropic retained contractual prohibitions on the use of Claude for "lethal autonomous warfare" and for "mass surveillance of Americans." *Id.* at 12–13.

In the fall of 2025, Anthropic and the Department began negotiations to establish a direct contractual relationship and to expand the Department's use of Claude. As part of that negotiation, the Department asked Anthropic for contractual permission to deploy Claude for "all lawful uses." App. 14. Anthropic agreed to substantially relax the prior use restrictions, but it continued to insist that Claude not be used for lethal autonomous warfare or mass surveillance of Americans, which it describes as "two critical exceptions." *Id.* Negotiations over these restrictions extended for a few months and eventually stalled.

C

The contractual dispute between Anthropic and the Department came to a head in early 2026.

On January 9, Secretary of War Pete Hegseth set forth a comprehensive "Artificial Intelligence Strategy for the Department." App. 202. According to the Secretary, "AI-enabled warfare and AI-enabled capability development will re-define the character of military affairs over the next decade." *Id.* The Secretary discerned a "race" between the United States and its "adversaries" to integrate AI technology into military capabilities. *Id.* The Secretary thus "direct[ed] the Department of War to accelerate America's Military AI Dominance by becoming an 'AI-first' warfighting force across all

components, from front to back." *Id.* As part of that directive, the Secretary stated that the Department must use AI models "free from usage policy constraints that may limit lawful military applications," and he directed the Department to "incorporate standard 'any lawful use' language into" contracts procuring AI services. *Id.* at 206.

Around the same time, an Anthropic executive "questioned the propriety" of a contractor's use of Claude "for a sensitive military operation abroad." App. 181. The Department believes that the governing usage policy "clearly permitted" the engagement at issue. *Id.* at 183. Nonetheless, the question "led to alarm by the DoW and the prime contractor who provides Anthropic software, and raised material doubts" about whether the software would perform as the Department was expecting. *Id.* at 228. The Department does not elaborate on the specific military operation at issue. However, Anthropic put into the record media reports stating that "Anthropic had raised concerns with Palantir about the role [Anthropic's] technologies played" in the January 3 "military operation to capture Venezuela's president, Nicolás Maduro." *Id.* at 171.

Finally, the Department learned of another instance when Anthropic's model training caused Claude to refuse to respond to queries from a government agency. Specifically, Claude refused to respond to queries from the Centers for Disease Control and Prevention (CDC) regarding sensitive research on preventing the spread of infectious disease. App. 212–13.

On February 24, Amodei met with Secretary Hegseth to discuss the standoff. The Secretary praised Claude's capabilities but demanded that Anthropic accede to an "all lawful uses" contractual term by February 27. App. 26.

On February 26, Anthropic refused, and Amodei released a statement explaining the company's decision to maintain the

two contested use restrictions. He stated that "mass domestic surveillance," although legal, was "incompatible with democratic values" and presented "serious, novel risks to our fundamental liberties." App. 146. He further stated that "*fully* autonomous weapons (those that take humans out of the loop entirely and automate selecting and engaging targets) may prove critical for our national defense" in the future, but that AI technology was not yet "reliable enough" to currently power such weapons. *Id.* at 147. Amodei recognized that "[i]t is the Department's prerogative to select contractors most aligned with their vision" for appropriate AI uses, and he pledged a "smooth transition to another provider" "[s]hould the Department choose to offboard Anthropic." *Id.*

One day later, President Trump and Secretary Hegseth denounced Anthropic's decision on social media, and the Secretary began the process of removing Claude from the Department's supply chain. App. 77, 153.

D

On March 3, 2026, Secretary Hegseth made a formal determination to take procurement actions against Anthropic under the Supply Chain Security Act. First, the Secretary determined that use of Claude in Department systems "presents a significant supply chain risk" and that removing Claude from them was "necessary to protect national security" by reducing that risk. App. 177. Next, he determined that no "less intrusive measures" for reducing the risk were "reasonably available." *Id.* Finally, he determined that an "urgent national security interest" required immediate action. *Id.* The determination rested on a recommendation from senior agency officials, which in turn rested on a memorandum from Emil Michael, the Department's Under Secretary for Research and Engineering. Among other things, Michael cited Anthropic's refusal to allow

all lawful uses of Claude, its ability to "alter system guardrails and model weights" governing how Claude responds to user prompts, and its questioning the Department's use of Claude in a sensitive military mission abroad. *Id.* at 181–83.

Secretary Hegseth immediately notified Anthropic of his determination, in a letter dated March 3 and emailed to Anthropic on March 4. App. 72, 243. The notice stated that the determination was effective immediately, and it gave Anthropic an opportunity to seek reconsideration within 30 days. *Id.* at 73.

The Department immediately began implementing the Secretary's determination. On March 6, its Chief Information Officer issued a Department-wide memorandum ordering the removal of Anthropic products from the Department's systems "as soon as practical," and in any event within 180 days. App. 80. The memo further prohibited contractors from using Anthropic products in their work for the Department. *Id.* The Department quickly moved to expand its contractual relationship with OpenAI, another company that provides AI services. *Id.* at 220. In response, Amodei wrote to Anthropic employees to express his view that the Department, OpenAI, and Palantir had not established adequate safety protocols for the use of AI. *Id.* at 220–23.

On March 9, Anthropic filed a petition for review of its exclusion from the Department's supply chain. Anthropic also moved for a stay pending review. In litigating the stay motion, the parties filed various affidavits and other evidentiary materials with this Court.

On March 19, the Department provided Anthropic with a supplemental notice, which included copies of the determination itself, the recommendation from agency officials, the memorandum by Under Secretary Michael, and a

statement of the exact scope of the covered procurement actions. App. 224. The supplemental notice restarted the 30-day deadline for Anthropic to submit any "information or arguments in opposition to this notice." *Id.*

On April 8, this Court denied a stay but expedited review on the merits. *Anthropic PBC v. U.S. Dep't of War*, No. 26-1049, 2026 WL 1042493 (D.C. Cir. Apr. 8, 2026) (per curiam). We ordered the parties to brief the question of our jurisdiction over the petition for review and to provide further information on how Anthropic could affect the functioning of its models before or after their delivery to the Department.

On April 17, Anthropic asked the Department to "rescind" its exclusion from the supply chain. *See* Letter from Counsel for Resp'ts to Clerk of Ct. at 6–7, *Anthropic PBC v. U.S. Dep't of War*, No. 26-1049 (D.C. Cir. filed May 12, 2026). There and in the appendix here, Anthropic tendered a supplemental declaration addressed to our factual question. App. 272–97. The Department has also tendered its own supplemental declaration, styled as one "in support of the administrative record." *Id.* at 408 (cleaned up); *see id.* at 408–22.

On June 3, the Secretary issued a decision denying reconsideration of his March 3 determination. Suppl. Br. for Resp'ts, Add. 1. The Secretary clarified that his determination did not rest on the premise that Anthropic could control any version of Claude after its delivery to contractors for deployment on the Department's classified systems. *See id.*

After oral argument in this Court, we ordered supplemental briefing on (1) whether Anthropic's April 17 filing divested this Court of jurisdiction and (2) the impact of the Secretary's June 3 order. Following that briefing, the matter is now ripe for decision.

III

We agree with the parties that we have jurisdiction to review the covered procurement actions taken against Anthropic under the Supply Chain Security Act.

To begin, we conclude that our jurisdiction was secure when Anthropic filed its petition for review on March 9. The Supply Chain Security Act gives this Court jurisdiction to review any "covered procurement action under section 4713," so long as the aggrieved party files the petition within 60 days of receiving notice of the covered action. 41 U.S.C. § 1327(b)(1). Here, the Secretary invoked section 4713 to remove Claude from the Department's supply chain. Anthropic received notice of the Secretary's action on March 4, and it promptly sought review five days later. That sufficed to establish jurisdiction under section 1327(b)(1).

We recognize that the statute distinguishes between the "covered procurement action[s]" authorized by section 4713(a), which are reviewable, and the antecedent written "determination" that use of the authority to take covered procurement actions is necessary, as separately required by section 4713(b). But the statute does not prohibit the government from making the written determination and taking the covered procurement actions at the same time. Here, the Department has done both: Its notice to Anthropic stressed that the determination had become "effective immediately." App. 73. And within three days, the Department had begun implementing the determination with an agency-wide order to "remove" Claude "from all DoW systems and networks … as soon as practical." *Id.* at 80. So, Anthropic was notified of "covered procurement action[s]" and timely sought review of them. 41 U.S.C. § 1327(b).

A distinct jurisdictional question arose after Anthropic asked the Department to rescind the covered procurement actions on April 17. Under statutory schemes limiting judicial review to final agency action, the filing of a motion for reconsideration renders incurably premature a previously filed petition for review. *See*, *e.g.*, *Nat'l Ass'n of Immigr. Judges v. FLRA*, 77 F.4th 1132, 1136–38 (D.C. Cir. 2023) (per curiam). Anthropic contends that this incurably-premature doctrine does not apply to judicial review under the Supply Chain Security Act because the statute contains no finality requirement. For its part, the Department contends that the doctrine does apply but that it is waivable and was waived in this case.

Subsequent developments make it unnecessary for us to resolve these questions. The Secretary denied Anthropic's request for rescission or reconsideration on June 3; Anthropic filed a petition for review of the March 3 and June 3 decisions on June 17; and we consolidated the two cases on June 24. We have jurisdiction over at least one of them. On the one hand, if the incurably-premature doctrine does not apply to judicial review under the Supply Chain Security Act, then Anthropic's request for rescission did not imperil our jurisdiction over its earlier-filed petition for review. On the other hand, if the incurably-premature doctrine does apply here and is not waivable, then the second petition for review cinched up our jurisdiction: Where the doctrine applies, the filing of a motion for reconsideration "tolls the period for judicial review of the original order, which can therefore be appealed to the courts *directly* after the petition for reconsideration is denied." *ICC v. B'hood of Locomotive Eng'rs*, 482 U.S. 270, 279 (1987). One way or the other, our jurisdiction is secure.

IV

Anthropic challenges each of the Secretary's three key determinations under the Supply Chain Security Act, namely that (A) removing Claude from the Department's supply chain was "necessary to protect national security by reducing supply chain risk," 41 U.S.C. § 4713(b)(3)(A); (B) "less intrusive measures" for reducing that risk were "not reasonably available," *id.* § 4713(b)(3)(B); and (C) an "urgent national security interest" required "the immediate exercise" of the removal authority, *id.* § 4713(c).

These claims are governed by the standards of review set forth in the Supply Chain Security Act, which requires us to "hold unlawful" covered procurement actions under section 4713 that we find to be "arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law." 41 U.S.C. § 1327(b)(2), (2)(A). These standards closely track those in the Administrative Procedure Act. *See* 5 U.S.C. § 706(2).

On review for arbitrariness, we require only that agency action be "reasonable and reasonably explained." *FCC v. Prometheus Radio Project*, 592 U.S. 414, 423 (2021). Judicial review under this standard is "deferential," and this Court may not "substitute its own policy judgment for that of the agency." *Id.* We will uphold even a "decision of less than ideal clarity," so long as the agency's basic rationale "may reasonably be discerned." *FCC v. Fox Television Stations, Inc.*, 556 U.S. 502, 513–14 (2009) (quoting *Bowman Transp., Inc. v. Arkansas-Best Freight Sys., Inc.*, 419 U.S. 281, 286 (1974)).

We review *de novo* agency determinations on purely legal questions of statutory construction. *Loper Bright Enters. v. Raimondo*, 603 U.S. 369, 412 (2024). In construing statutes implicating national security, we resolve ambiguities against intruding on the national-security determinations of the

Executive Branch. *See*, *e.g.*, *Dep't of Navy v. Egan*, 484 U.S. 518, 530 (1988); *Hikvision USA, Inc. v. FCC*, 97 F.4th 938, 948 (D.C. Cir. 2024); *Fed. Express Corp. v. U.S. Dep't of Com.*, 39 F.4th 756, 769 (D.C. Cir. 2022).

Where the application of a statutory term turns primarily on factual determinations, we review the agency's assessment deferentially. *Seven County Infrastructure Coal. v. Eagle County*, 605 U.S. 168, 181 (2025). And where a factual question turns on assessments of national security, we give the agency more deference, even if constitutional claims are at issue. *See*, *e.g.*, *Trump v. Hawaii*, 585 U.S. 667, 704 (2018) ("our inquiry into matters of … national security is highly constrained"); *Holder v. Humanitarian L. Project*, 561 U.S. 1, 33 (2010) (*HLP*) (in national-security cases, "evaluation of the facts by the Executive, like Congress's assessment, is entitled to deference").

A

The Secretary reasonably concluded that removing Anthropic from the Department's supply chain was necessary to protect national security by reducing supply chain risk to the Department's information systems. Specifically, the Secretary credited a joint recommendation from two senior Department officials that Claude might be "subject to manipulation" by Anthropic "in such a manner as to inhibit the DoW's use thereof." App. 178. Likewise, he credited Under Secretary Michael's conclusion that Anthropic might "manipulat[e]" the "design, integrity, and operation" of the Department's Claude models, potentially causing "critical defense system[s] failing to engage" as intended by the Department. *Id.* at 182.

1

The record in this case amply supports the Secretary's conclusion. To begin, it is undisputed that Anthropic can and does control how Claude responds—or fails to respond—to user prompts. Anthropic's Chief Science Officer explained how the company "seek[s] to embed safety considerations directly into the model itself." App. 8. Its CEO explained how such training gives the model an "identity, character, values, and personality" of its own, tethered to a "constitution" developed to impose "high-level principles and values" on Claude itself. *Id.* at 93–94. And the head of its public-sector business explained: "Model training is the primary mechanism through which Anthropic can influence the behavior of models used by the Department." *Id.* at 276. Anthropic disclaims any ability to access or alter a model that has already been delivered to the Department, *see id.*, despite the "technical measures" that it uses to police compliance with usage restrictions by private customers, *id.* at 8. Nonetheless, extant models reflect Claude's "[c]onstitutional" training. *Id.* at 274–75. Moreover, Anthropic may encode additional restrictions each time it delivers any "new version of the model" to Department contractors. *Id.* at 276. Finally, it is undisputed that such model restrictions are vitally important to Anthropic, which describes them as lying "at the core of [its] mission." *Id.* at 2.

The record also indicates that Anthropic's model training is effective in enforcing usage restrictions and that, as a result, Claude has refused to answer legitimate queries from government users. Anthropic itself explained how early, commercially available versions of Claude frustrated Department and intelligence-community users by refusing prompts to evaluate classified materials. App. 255. Likewise, as Under Secretary Michael explained, the Department learned in 2025 that Claude had refused to process CDC prompts to

support research to prevent the spread of infectious diseases. *Id.* at 212. Anthropic responds that these glitches reflected safety features appropriately built into models sold to private companies and were resolved after Anthropic engineers worked with the relevant government stakeholders. *Id.* at 255–56, 261–62. Perhaps so, but the point here is not that these model or usage restrictions were arbitrary; instead, it is that Anthropic's model training does effectively enforce contractual usage restrictions.

Finally, the record reveals a recent, serious dispute about the scope of the contractual prohibitions on lethal autonomous warfare and mass domestic surveillance. Under Secretary Michael describes the incident in general but striking terms:

> [O]ne of Anthropic's executives questioned the propriety of the potential use of their software for a sensitive military operation abroad despite that use being permitted under the existing Terms of Service. This led to alarm by the DoW and the prime contractor who provides Anthropic software, and raised material doubts as to whether they would cause their software to stop working or cause some other disastrous action that would put our warfighters['] lives in danger.

App. 181. Anthropic does not say much about this incident, except to suggest that it reflected a misunderstanding. *Id.* at 236–37. But regardless, Anthropic has made clear that it views the contractual prohibition on mass domestic surveillance as urgent to support "democratic values," *id.* at 146, and the contractual prohibition on lethal autonomous warfare as urgent to avoid "put[ting] America's warfighters and civilians at risk" of a catastrophic AI mistake, *id.* at 147. For its part, the Department has made clear that it views an "any lawful use" authorization to be critical to its "AI-first" strategic plan. *Id.* at

202, 206. With such diametrically opposed positions and with contractual limitations that are hardly self-defining, the prospect for disputes is apparent.

Because Anthropic was willing and able to enforce contractual restrictions through model training, the Department reasonably worried that "critical defense system[s]" supported by Claude might "fail[] to engage" as the Department would expect. *Id.* at 182 (Michael memorandum). Of course, we do not know exactly what happened in the incident described by Michael as a near-disaster and by Amodei as a misunderstanding. But Anthropic's suggestion that the incident may have arisen during a shock-and-awe, kinetic operation to capture a foreign head of state abroad simply underscores the fraught nature of its dispute with the Department—and the Department's need for certainty that its AI systems will perform as expected.

In sum, the Department reasonably feared that Anthropic might manipulate Claude's design to prevent it from performing national-security functions that the Department deems contractually authorized and necessary.

2

Anthropic offers two primary responses, one factual and one legal. Neither is persuasive.

a

On the facts, Anthropic stresses that it cannot control or even monitor the operation of any model once the model is delivered to a Department contractor for use on a classified system. App. 274–78. In sum, Anthropic says it has no "back door or remote 'kill switch.'" *Id.* at 258. So, it reasons, the Department can test any new model delivered by Anthropic to

contractors—before integrating it into the Department's information systems—to see if the new model performs up to the Department's expectations. *Id.* at 282–92. Specifically, Anthropic says that the Department can "test the model" to see if it will "refuse tasks the Department deem[s] appropriate to its lawful mission, or … override the Department's judgment that an activity is permissible." *Id.* at 285. And if the Department is not satisfied with the new model, it can simply "decline to approve" the model. *Id.* at 286. This would leave the Department free to continue using an older, previously accepted model, which "does not degrade or change on its own." *Id.* at 287–88.

These responses do not assuage the Department's concerns. To begin with, the Department has good reasons not to view advance testing as a panacea. As noted above, the contested use restrictions are hardly self-defining, and there are any number of possible scenarios involving, say, greater or lesser degrees of human involvement in targeting decisions during ongoing military operations. Moreover, as Under Secretary Michael explained, Claude employs technology that is "opaque" to its users, App. 182, partly because "Anthropic's unique building processes are considered proprietary intellectual property" and partly because its models "have weights or parameters that number approximately 5 to 10 trillion per model," *id.* at 410. All of this makes "rigorous analysis or auditing of its output mathematically impossible." *Id.* Indeed, Anthropic itself acknowledges "some legitimacy to DoW's concern about the opacity of these systems generally." *Id.* at 261. Finally, Anthropic acknowledges that Claude might respond differently to similar requests depending on their exact wording. *Id.* at 289. So, while one Department official might secure an advance commitment from Claude to perform a contemplated military function, a second official, using slightly different wording, might later be met with a refusal.

In any event, even if testing could reliably establish in advance the range of situations where a new model might decline to perform some lawful but contractually prohibited function, the Department could hardly be satisfied with the option of simply refusing the upgrade. Anthropic "continually develop[s] and release[s] increasingly capable versions of Claude." App. 6. In 2025, it released three new, increasingly powerful versions of Claude Gov. *Id.* at 256–57. Since then, it has released several more versions of Claude. *See* Anthropic, *Models Overview*, https://perma.cc/BR4B-SRYJ. The Secretary has noted the "unprecedented velocity in the evolution of the frontier AI models," which "are becoming smarter and more robust every day." App. 205. Amodei likewise acknowledges "a smooth, unyielding increase in AI's cognitive capabilities" over the last few years, and he predicts even greater breakthroughs on the horizon, such as a "point where the current generation of AI autonomously builds the next." *Id.* at 85–86. More ominously, he also foresees—as perhaps do our adversaries—that a "swarm of millions or billions of fully automated armed drones, locally controlled by powerful AI and strategically coordinated across the world by an even more powerful AI, could be an unbeatable army." *Id.* at 106. Quite obviously, the Department cannot utilize AI systems that remain trapped in amber.

Finally, Anthropic claims that Under Secretary Michael's memorandum contained a discrete factual error in suggesting that Anthropic could modify the behavior of models already delivered to Department contractors. In the one sentence at issue, Michael expressed concern that Anthropic might have retained the ability to "disable its technology … in advance *or in the middle of* ongoing warfighting operations." App. 183 (emphasis added). Anthropic has since clarified that once a model is delivered for use on the Department's classified systems, Anthropic cannot "access it, alter it, or shut it down."

Id. at 276. But as explained above, Anthropic can and does program Claude's behavior with each new model it delivers to Department contractors. So the basic thrust of Michael's analysis remains intact: Anthropic retains the ability to "alter system guardrails and model weights" over time, and it can use that ability to prevent Claude from "engag[ing]" in specific operations that *it* deems to reflect contractually unauthorized uses. *Id.* at 182. Moreover, the Secretary's order denying reconsideration removes any doubt on whether this specific objection matters. In that order, the Secretary clarified that his determination "did not depend" on any particular understanding of Anthropic's "real-time technical access to or control" of Claude "post-deployment on the Department's covered systems." Suppl. Br. for Resp'ts, Add. 1. Nor, for that matter, did his denial of reconsideration depend on any such understanding. *Id.* at 1–2.

b

Alternatively, Anthropic contends that the Secretary's concerns about what it might do to restrict Claude's functionality, even if factually supported and reasonably explained, do not amount to a covered "supply chain risk." The statute defines that term to mean:

> the risk that any person may sabotage, maliciously introduce unwanted function, extract data, or otherwise *manipulate* the design, integrity, manufacturing, production, distribution, installation, operation, maintenance, disposition, or retirement of covered articles so as to surveil, *deny*, disrupt, or otherwise manipulate the function, use, or operation of the covered articles or information stored or transmitted on the covered articles.

41 U.S.C. § 4713(k)(6) (emphases added).

At first glance, Anthropic's argument runs headlong into seemingly clear statutory terms. To *manipulate* is "[t]o move, arrange, operate, or control by the hands or another body part or by mechanical means, especially in a skillful manner." *Manipulate*, American Heritage Dictionary, https://perma.cc/B2FA-CM84. And to *deny* is simply "[t]o decline to grant or allow." *Deny*, American Heritage Dictionary, https://perma.cc/3WB4-GF3M. Based on undisputed record evidence, there is not only a "risk"—but a certainty—that Anthropic will so *manipulate* the "design" or "operation" of Claude to *deny* it the "function" of conducting lethal autonomous warfare or mass domestic surveillance.

Anthropic seeks to impose a narrowing construction on these statutory terms. It notes that one of the verbs in the first relevant string (*sabotage*) connotes intentionally hostile acts and that another of the verbs (*introduce*) is limited by an adverb (*maliciously*) that likewise connotes a bad motive. Anthropic therefore asks us to impose the same limitation—of intentional hostility or bad motive—on the entire statutory definition. Anthropic spends less than two pages of its opening brief on this interpretive argument, Br. for Pet'r at 39–41, but the dissent amplifies it considerably. To narrow the statutory definition based on contextual considerations, the dissent invokes the associated-words canon, the *ejusdem generis* canon, the series-qualifier canon, and the surplusage canon. It also quotes repeatedly from the dissent in *United States v. Fischer*, 64 F.4th 329, 363–83 (D.C. Cir. 2023) (Katsas, J., dissenting), which urged a narrow contextual reading of a different statute, and from the Supreme Court decision adopting that reading, *Fischer v. United States*, 603 U.S. 480 (2024). The dissent's arguments here have some force, but we are ultimately unpersuaded.

Start with the first string of relevant verbs—*sabotage*, *introduce*, *extract*, and *manipulate*. This string does not uniformly connote bad motive; *sabotage* does, but *extract* and *manipulate* do not, and *introduce* does only as modified by the adverb *maliciously*. Moreover, these verbs do not appear in a tight parallel list; one stands on its own (*sabotage*), while the other three take distinct direct objects (*introduce* goes with *function*, *extract* goes with *data*, and *manipulate* goes with a long string of nouns including *design* and *operation*). And the specific verb-object phrase invoked by the Department (to *manipulate* the *design* or *operation* of an information-technology product) perfectly describes Anthropic's constitutional training of Claude. Neither the associated-words nor the *ejusdem generis* canons support artificially narrowing *manipulate* to require a bad motive, because the verbs preceding *manipulate* do not uniformly connote acts done with bad motive. *See* A. Scalia & B. Garner, *Reading Law* 196 (2012) ("For the associated-words canon to apply, the terms must be conjoined in such a way as to indicate that they have some quality in common."); *id.* at 199 (*ejusdem generis* canon) ("When the initial terms all belong to an obvious and readily identifiable genus, one presumes that the speaker or writer has that category in mind for the entire passage."). And the grammatical complexity of the entire phrase—with intervening direct objects, adverbs, and adjectives—cuts further against application of these canons by tending to weaken the required "association" among the listed verbs. *See id.* at 197.

To overcome these substantial problems, the dissent invokes the series-qualifier canon to contend that the adverb *maliciously*, as used in the phrase "maliciously introduce unwanted function," modifies not only the immediately following verb *introduce*, but also the next verb *extract*. *Post*, at 5. On that reading, *sabotage*, *maliciously introduce*, and *maliciously extract* all would connote acts done with bad

motive, which would lay more groundwork for application of the associated-words or *ejusdem generis* canons. *Post*, at 4–6. But the series-qualifier canon does not fit. Under that canon, "[w]hen there is a straightforward, parallel construction that involves all nouns or verbs in a series, a prepositive or postpositive modifier normally applies to the entire series." A. Scalia & B. Garner, *supra*, at 147; *see Facebook, Inc. v. Duguid*, 592 U.S. 395, 402 (2021). In other words, the canon "generally applies when a modifier precedes or follows a list, not when the modifier appears in the middle." *Wong v. Minn. Dep't of Hum. Servs.*, 820 F.3d 922, 928 (8th Cir. 2016). In the phrase "Constitution, statutes, or treaties of the United States," the adjectival *of the United States* plainly modifies all three nouns. *See Lockhart v. United States*, 577 U.S. 347, 363–69 & nn.1–2 (2016) (Kagan, J., dissenting). So does the adjectival *United States* in the phrase "United States Constitution, statutes, or treaties." But not so for *of the United States* in the phrase "Constitution, statutes of the United States, or treaties." Here, the word *maliciously* modifies the second of four verbs in the sequence, so there is scant basis for projecting it backwards to *sabotage*, or forwards to *extract* or *manipulate*. Moreover, as explained above, the list of verbs is neither straightforward nor parallel; instead, it has "unexpected internal modifiers or structure" as well as "varied syntax," which further suggests that *maliciously* modifies only its closest verb: *introduce*. *Lockhart*, 577 U.S. at 352 (majority opinion); *see* A. Scalia & B. Garner, *supra*, at 152 ("When the syntax involves something other than a parallel series of nouns or verbs, a prepositive or postpositive modifier normally applies only to the nearest reasonable referent.").

The dissent answers with this hypothetical library rule: "Do not shout, loudly talk on the phone, play music, or otherwise disturb others." *Post*, at 5. The dissent posits that *loudly* modifies "play music" as well as "talk on the phone," in

the middle of a verb sequence roughly parallel to the one at issue here. *Id.* The dissent may be correct that only loud music violates the library rule, but not because adverbs in the middle of a verb string tend to jump forwards or backwards as if the series-qualifier canon applied. Instead, the dissent's interpretation of its hypothetical rule sounds plausible because the words preceding a residual *otherwise* clause normally provide "examples" of the words that follow, *Fischer*, 603 U.S. at 487; *see Begay v. United States*, 553 U.S. 137, 144 (2008), and it is hard to understand playing inaudible music with headphones as an example of something that might *disturb* other library patrons. In this respect, the statutory definition here is different; read the phrase "extract data" to cover non-malicious actions consistent with its ordinary meaning, and it still provides an easily recognizable example of manipulating the design or operation of an information-technology system. So there is no reason to awkwardly project forward the adverb *maliciously* from *introduce* to *extract*.

The second verb string in the definition strengthens the case for a plain-meaning interpretation of *manipulate*. That tight parallel string ("surveil, deny, disrupt, or otherwise manipulate") does not convey any overarching connotation of bad motive. In this string, one of the three verbs before *manipulate* has no such connotation (*deny*), and the other two are at worst ambiguous on this point (*surveil* and *disrupt*). So the second string of verbs does not support application of the associated-words or *ejusdem generis* canons. The dissent all but recognizes as much, but proposes reading the two strings together. *Post*, at 5–6. Fair enough, but doing so weakens its case: The dominance of neutral or ambiguous verbs in the second string suggests that *manipulate* does *not* bear a contextually narrowed meaning requiring malice, and that in turn suggests that *manipulate* in the first string also does not.

See A. Scalia & B. Garner, *supra*, at 170–73 (presumption of consistent usage).

Finally, the statute is ultimately addressed to concerns about the proper "function, use, or operation" of the covered product, which focuses on the effect of a manipulation, not the intent behind it. And denied or disrupted function in the government's information-technology systems might well implicate national security. So, on balance, we see little reason to depart from the most ordinary meaning of *manipulate*—to skillfully arrange, operate, or control. We have no reason to doubt that Anthropic manipulates Claude's function, use and operation with noble intentions, whether a principled commitment to personal privacy or a genuine concern about AI safety. But at least as applied here, the statutory definition of a "supply chain risk" turns on *what* Anthropic does, not *why* Anthropic does it.

Broader statutory context reinforces our conclusion. Section 4713 imposes no criminal liability, a context in which we would strive to narrow the statute or impose a *mens rea* requirement on it. *See*, *e.g.*, *Fischer*, 603 U.S. at 496–97; *Elonis v. United States*, 575 U.S. 723, 734 (2015). Nor does it even impose any civil liability. To the contrary, section 4713 is a procurement statute addressed to what goods and services the Department may buy, and even Anthropic disclaims any challenge to the Department's "fundamental prerogative" to choose its contractors and subcontractors. Br. for Pet'r at 1. Moreover, section 4713 is not just any procurement statute, but one enabling the Executive Branch to mitigate national-security risks. As noted above, we construe ambiguities in such statutes *in favor* of the government, to avoid needlessly or dangerously constraining the national-security determinations of the Executive Branch. *See*, *e.g.*, *Lee v. Garland*, 120 F.4th 880, 888 (D.C. Cir. 2024) ("generally applicable statutes

should not be applied to impinge on" executive-branch control over security clearances "absent some clear statement by Congress"); *Fed. Express Corp.*, 39 F.4th at 769 ("courts accord special deference to an agency construction of a statute 'in the areas of foreign policy and national security'" (quoting *Haig v. Agee*, 453 U.S. 280, 291 (1981))); *Changji Esquel Textile Co. v. Raimondo*, 40 F.4th 716, 723 (D.C. Cir. 2022) ("Our interpretive approach must also account for the substantial deference due to the Executive Branch in this context."); *Al-Bihani v. Obama*, 619 F.3d 1, 39–40 (D.C. Cir. 2010) (Kavanaugh, J., concurring in the denial of rehearing en banc) (explaining courts' "traditional deference in interpreting national security statutes").

Fischer itself does not help Anthropic any more than the canons invoked by the dissent. In that case, the predicates for invoking the associated-words and *ejusdem generis* canons were satisfied because the examples before the residual *otherwise* clause all clearly involved the impairment of evidence—the contextual limitation imposed by the Court on the residual clause. *See* 603 U.S. at 487, 489–90. Moreover, a non-contextual reading of the residual clause in *Fischer* would have produced massive surplusage—rendering superfluous all of the reticulated examples preceding that clause, 15 of the 21 offenses in the statutory code section at issue, and much of the entire corpus of obstruction offenses. *See id.* at 492–94; 64 F.4th at 371–73 (Katsas, J., dissenting). Nothing like that is present here, where the residual clause is itself longer and more reticulated than the three brief examples preceding it. Furthermore, a literalist construction of the residual clause in *Fischer* would have implausibly extended the prohibition at issue to constitutionally protected activity such as advocacy, lobbying, and protest, triggering concerns of constitutional avoidance. *See* 603 U.S. at 496; 64 F.4th at 378–79 (Katsas, J., dissenting). Here, as explained above, non-maliciously

disrupting the function or denying the use of the government's information systems can plausibly raise national-security concerns and does not implicate constitutionally protected conduct. Finally, *Fischer* involved a criminal prohibition, triggering the obligation to read the statute narrowly if possible, whereas this case involves a national-security procurement statute, triggering the opposite interpretive presumption.

Anthropic also invokes the recent decision in *Anthropic PBC v. U.S. Dep't of War*, No. 26-cv-01996 (N.D. Cal. Aug. 27, 2026), which set aside the Department's designation of Anthropic as a "supply chain risk" under 10 U.S.C. § 3252. But the definition of that term under section 3252 is much narrower than the corresponding definition under 41 U.S.C. § 4713. Under section 3252, the term "supply chain risk" means "the risk that an adversary may sabotage, maliciously introduce unwanted function, or otherwise subvert the design, integrity, manufacturing, production, distribution, installation, operation, or maintenance of a covered system so as to surveil, deny, disrupt, or otherwise degrade the function, use, or operation of such system." 10 U.S.C. § 3252(d)(4). We have no quarrel with the Northern District's conclusion that use of the critical noun *adversary*, combined with the sinister connotation fairly pervading the string of *sabotage*, *maliciously introduce*, and *otherwise subvert*, indicate that bad motive is required to support a designation under section 3252. Likewise, we have no quarrel with the Northern District's conclusion that Anthropic has acted with no such bad motive in its dealings with the Department. But as explained at length above, no such bad motive is required to support a designation under the much broader definition set forth in section 4713.[1]

[1] Anthropic contends that the Northern District's decision is preclusive as well as persuasive. But because the Department's designation authority is much broader under section 4713 than it is

Two final points. First, for the reasons discussed above, we reject Anthropic's attempt to engraft onto the statutory definition an overarching requirement of acting surreptitiously, just because two of the seven verbs in the strings (*sabotage* and *surveil*) have that connotation. Second, we reject Anthropic's attempt to glean from the legislative history a focus on "foreign companies working at the behest of foreign states." Br. for Pet'r at 40. Whatever paradigmatic examples individual members of Congress may have had in mind, the statutory definition is not limited to "adversar[ies]," 10 U.S.C. § 3252(d)(4), and instead covers "any person," which cannot refer only to foreign entities, 41 U.S.C. § 4713(k)(6). Likewise, we reject the dissent's use of legislative history, *post*, at 7–8, to glean the ordinary public meaning of the seven verbs in the definition.

In sum, we conclude that the Secretary's concern about Anthropic disabling Claude from performing lawful actions

---

under section 3252, the issues flagged by Anthropic are not the same in both cases. So, for example, the Northern District's determination that the section 3252 designation was arbitrary does not control our determination whether the section 4713 designation was arbitrary. Likewise, the Northern District's determination of exigency under section 3252 does not control our determination of exigency under section 4713. In any event, Congress gave this Court exclusive jurisdiction to review procurement actions taken pursuant to section 4713 designations, *see* 41 U.S.C. § 1327(b)(1), and it specifically barred other courts from reviewing any other "action taken under" section 4713, *see id.* § 1327(a). That strict "allocation of jurisdiction" to this Court makes it inappropriate to constrain our review based on the Northern District's judgment. Restatement (Second) of Judgments § 28 (1982); *see Shaw v. State of Cal. Dep't of Alcoholic Beverage Control*, 788 F.2d 600, 607–09 (9th Cir. 1986); *Lyons v. Westinghouse Elec. Corp.*, 222 F.2d 184, 188–89 (2d Cir. 1955) (L. Hand, J.).

requested by the Department qualifies as a "supply chain risk" within the meaning of section 4713.

B

The Secretary also concluded that less intrusive measures were not reasonably available to reduce the supply-chain risk posed by Anthropic. We have already explained that the Department's concerns are serious and fall within the statutory definition. Anthropic all but acknowledges that such concerns would justify the termination of its subcontracts with the Department. Even so, Anthropic objects that the Department's invocation of section 4713, as opposed to more conventional procurement authorities, had the additional effect of "branding" the company as a national-security threat, which assertedly harmed its reputation. Reply Br. at 21–22. One may fairly question whether Anthropic has suffered any such harm, as rapid increases in the company's valuation since the Secretary's section 4713 determination reportedly have made it one of the most valuable business concerns in the world. *See*, *e.g.*, Clark, *Anthropic Was Behind. Now It's the AI Boom's Front-Runner*, Wall St. J. (May 13, 2026) ("Anthropic has received investment offers in recent months valuing it at more than $900 billion"). Regardless, Anthropic has not explained why it would suffer any less of a stigma if the Department had articulated the very same national-security concerns to announce the elimination of Claude from its supply chain through more conventional contract-termination authorities.

Alternatively, Anthropic suggests in one sentence that the Department should have "[n]arrow[ed] any restriction to the subset of systems plausibly involving lethal autonomous warfare or domestic mass surveillance." Br. for Pet'r at 48. This fleeting statement did not preserve the point. *Schneider v. Kissinger*, 412 F.3d 190, 200 n.1 (D.C. Cir. 2005). In any

event, the suggestion overlooks the fact that, when an AI model is "layered into other applications," the model can "limit the functionality of that larger system." App. 214–15. Once the Secretary identified the risk that Claude posed, he sought to avert it by making a clean break instead of courting the delay, expense, and uncertainty that would result from a granular inquiry into every possible use of Claude by each of the Department's prime contractors. The Secretary reasonably concluded that less intrusive measures were not reasonably available, and we cannot override that fact-based, national-security assessment.

C

Anthropic contends that the Secretary impermissibly found that an "urgent national security interest require[d] the immediate exercise" of his authority to exclude Anthropic from the Department's supply chain. 41 U.S.C. § 4713(c). Without such an emergency, the Secretary would have had to provide Anthropic with advance notice and an opportunity to respond before making the exclusion. *See id.* § 4713(b). Anthropic faces strong headwinds in asking us to override the Secretary of War's assessment of national-security exigencies. But even assuming that this assessment was both reviewable and wrong, Anthropic cannot show any prejudice from the timing of the Secretary's notice.

1

"In administrative law, as in federal civil and criminal litigation, there is a harmless error rule." *Combat Veterans for Cong. Pol. Action Comm. v. FEC*, 795 F.3d 151, 156–57 (D.C. Cir. 2015) (quoting *Nat'l Ass'n of Home Builders v. Defs. of Wildlife*, 551 U.S. 644, 659–60 (2007)); *see also Shinseki v. Sanders*, 556 U.S. 396, 406 (2009) (harmless-error rules "ordinarily apply in civil cases"). The rule traces back to pre-

APA caselaw, where courts would require a prejudicial error in order to set aside agency action. *See Mkt. St. Ry. Co. v. R.R. Comm'n of State of Cal.*, 324 U.S. 548, 562 (1945). The APA then codified the rule, directing that reviewing courts take "due account … of the rule of prejudicial error." 5 U.S.C. § 706. As early commentators observed, the APA merely "sum[med] up … the 'harmless error' rule applied by the courts in the review of lower court decisions as well as of administrative bodies, namely, that errors which have no substantial bearing on the ultimate rights of the parties will be disregarded." Att'y Gen.'s Manual on the Admin. Proc. Act 110 (1947) (APA's prejudicial-error rule "appears to restate existing law"). Moreover, some form of the prejudicial-error rule is ubiquitous in civil cases, criminal cases, and appeals. *See* 28 U.S.C. § 2111 (in any appeal, courts must apply a "[h]armless error" rule, disregarding "errors or defects which do not affect the substantial rights of the parties"); Fed. R. Civ. P. 61 ("At every stage of the proceeding, the court must disregard all errors and defects that do not affect any party's substantial rights."); Fed. R. Crim. P. 52(a) ("Any error, defect, irregularity, or variance that does not affect substantial rights must be disregarded."). So "the 'harmless error' principle announced for our general jurisprudence by decision and statute … is applicable to the review of the decisions of administrative agencies"—even outside the APA itself. *Braniff Airways, Inc. v. Civil Aeronautics Bd.*, 379 F.2d 453, 465 (D.C. Cir. 1967).

We have repeatedly applied such prejudicial-error rules in judicial-review schemes that do not explicitly impose them. For example, in *Chai v. Dep't of State*, 466 F.3d 125 (D.C. Cir. 2006), we found harmless error in reviewing the designation of an entity as a foreign terrorist organization. *Id.* at 132–33. The governing judicial-review scheme there, set forth in the Antiterrorism and Effective Death Penalty Act (AEDPA), imposed APA-like standards of review but no express

requirement of prejudicial error. *See* 8 U.S.C. § 1189(c). Likewise, in *Saunders v. Kijakazi*, 6 F.4th 1 (D.C. Cir. 2021), we conducted harmless-error review (but found the errors there prejudicial) in a case governed by the judicial-review provisions of the Social Security Act, which also does not expressly require that courts consider harmless error. *Id.* at 4; *see* 42 U.S.C. § 405(g). These cases accord with a background principle that reviewing courts must not become "impregnable citadels of technicality." *Shinseki*, 556 U.S. at 407–08 (quoting *Kotteakos v. United States*, 328 U.S. 750, 759 (1946)). More generally, they confirm another background principle that "the law cares not for trifles." *Wisc. Dep't of Revenue v. William Wrigley, Jr., Co.*, 505 U.S. 214, 231 (1992).

Here, Anthropic failed to show any prejudice from not having received an advance opportunity to respond. By March 19, just over two weeks after the determination to exclude Claude from the Department's supply chain, the agency had provided Anthropic with notice of that action and the materials on which it was based. App. 72, 224. The Department also invited Anthropic to submit any "information or arguments in opposition to this notice." *Id.* at 224. Anthropic made such a submission, including a letter from counsel and the various evidentiary materials also submitted to this Court. The prejudice inquiry thus boils down to the question whether these materials, if provided sooner, may have changed the Department's decision. *See PDK Labs, Inc. v. DEA*, 362 F.3d 786, 799 (D.C. Cir. 2004) ("If the agency's mistake did not affect the outcome, if it did not prejudice the petitioner, it would be senseless to vacate and remand for reconsideration.").

We are confident that, just as these materials did not convince the Department to rescind the supply-chain exclusion in June, they would not have affected the Department's decision to initiate the exclusion in March. As the Secretary

explained on rehearing, much of the information submitted by Anthropic—such as the parties' contractual and negotiating history—was "already known to the Department." Suppl. Br. for Resp'ts, Add. 1. And much of the color and nuance in the materials—regarding just how deeply committed Anthropic is to its model training and use restrictions—reinforced the concerns expressed by Under Secretary Michael. The only new, potentially helpful piece of information provided by Anthropic was its explanation that, although the company uses "technical measures" to monitor ongoing use and prevent misuse by private customers, App. 8, it has no such ability with respect to models deployed on the Department's classified systems, *id.* at 287–88. But as the Secretary made clear in denying rehearing, his original decision did not rest on the premise that Anthropic was willing and able to monitor ongoing use, and prevent perceived misuse, by the Department once it had received the relevant models. Suppl. Br. for Resp'ts, Add. 1. In reviewing the Department's decision, we may consider this "amplified articulation" of its original rationale. *DHS v. Regents of the Univ. of Cal.*, 591 U.S. 1, 20 (2020); *see also Camp v. Pitts*, 411 U.S. 138, 143 (1973).

2

Anthropic's counterarguments are unpersuasive.

First, it contends that harmless-error review cannot apply in this case because the governing judicial-review provision here requires that this Court "shall hold unlawful" supply-chain designations "not in accord with procedures required by law." 41 U.S.C. § 1327(b)(2), (2)(E). But the AEDPA provision for judicial review of foreign terrorist organization designations uses equally mandatory language, in stating that the reviewing court "shall hold unlawful and set aside" FTO designations "not in accord with the procedures required by law." 8 U.S.C.

§ 1189(c)(3), (3)(E). Yet we denied review of a due-process claim on the ground that "the alleged errors were, in the particular circumstances of th[at] case, clearly rendered harmless." *Chai*, 466 F.3d at 132. Moreover, as explained above, the rule of prejudicial error is a ubiquitous feature of judicial-review schemes despite the equally ubiquitous presence of such "shall set aside" provisions.

Second, Anthropic contends that delayed notice here was prejudicial because "psychological and bureaucratic realities" might have locked the Department into its exclusion regardless of the strength of any later showing. *New Jersey, Dep't of Env't Prot. v. EPA*, 626 F.2d 1038, 1050 (D.C. Cir. 1980). In the case Anthropic invokes, we held that an agency's failure to follow notice-and-comment rulemaking procedures is prejudicial despite an opportunity to submit comments after promulgation of the final rule. *See id.* at 1049. In significant part, we reasoned that many prospective commenters may be unlikely to bother once a final rule was already in place. *See id.* Here, in contrast, the only question is whether an earlier submission *by Anthropic* would have made any difference, and we have no reason to think that the timing of its actual submission caused Anthropic to pull any punches. Moreover, Anthropic does not explain why its preferred remedy—holding the exclusion unlawful for lack of advance notice and an opportunity to be heard—would be any less susceptible to the same bureaucratic objection. We can hold unlawful the exclusion on procedural grounds, but we cannot prevent the Secretary from considering whether to reimpose it after reviewing whatever materials Anthropic may submit. Nor can we make the Secretary forget that, after considering Anthropic's full submission the last time around, he chose to maintain the exclusion. Under these circumstances, requiring another go-round would be pointless—the very type of futile gesture that the prejudicial-error rule avoids.

Third, Anthropic argues that reviewing for harmless error would eviscerate the advance-notice requirement because the agency could freely skip it, offer post-exclusion process, and then be insulated from judicial review. But even with a harmless-error rule in place, the advance-notice requirement still has teeth. Immediately after a suspect urgency determination, an aggrieved party could seek interim relief before any post-exclusion process had run its course. To demonstrate prejudice in that context, the petitioner would simply need to show that it could "mount a credible challenge" to the exclusion if afforded the opportunity. *Gerber v. Norton*, 294 F.3d 173, 184 (D.C. Cir. 2002) (quoting *Util. Solid Waste Activities Grp. v. EPA*, 236 F.3d 749, 755 (D.C. Cir. 2001)). But the analysis changes after a petitioner has received the post-designation process, which can cure the lack of earlier process. When a petitioner has *already* mounted a failed challenge—as in this case—it can no longer show prejudice from failing to receive an earlier opportunity to be heard.

For these reasons, we hold that Anthropic has suffered no prejudice from the Secretary's determination that an urgent national-security interest required Claude's immediate exclusion from the Department's supply chain.

V

Finally, Anthropic contends that its exclusion from the Department's supply chain violates the Fifth and First Amendments to the Constitution.

A

The Fifth Amendment prohibits the federal government from depriving any person of "life, liberty, or property, without due process of law." We assume that the Department's exclusion of Anthropic from its supply chain deprived the

company of protected liberty or property interests. Anthropic contends that the "due" process for such deprivations included a right to contest the exclusion *before* it became effective. As a general matter, due process requires notice and an opportunity to respond "before the government can constitutionally deprive a person of the protected liberty or property interest." *Nat'l Council of Resistance of Iran v. Dep't of State*, 251 F.3d 192, 205 (D.C. Cir. 2001) (*NCRI*). But "due process is flexible and calls for such procedural protections as the particular situation demands." *Id.* (quoting *Morrissey v. Brewer*, 408 U.S. 471, 481 (1972)). So, "where [the government] must act quickly, or where it would be impractical to provide predeprivation process, postdeprivation process satisfies the requirements of the Due Process Clause." *Zevallos v. Obama*, 793 F.3d 106, 116 (D.C. Cir. 2015) (quoting *Gilbert v. Homar*, 520 U.S. 924, 930 (1997)) (cleaned up). Post-deprivation process thus suffices where the government seeks to seize movable assets, *id.*; destroy unwholesome food, *N. Am. Cold Storage Co. v. City of Chicago*, 211 U.S. 306, 315 (1908); or terminate disability benefits, *Mathews v. Eldridge*, 424 U.S. 319, 349 (1976). Likewise, it suffices where "earlier notification would impinge upon the security and other foreign policy goals of the United States." *People's Mojahedin Org. of Iran v. Dep't of State*, 613 F.3d 220, 227 n.4 (D.C. Cir. 2010) (quoting *NCRI*, 251 F.3d at 208).

The need to move quickly was present here, as reflected in the Secretary's finding that immediate action was necessary. The parties agree that AI technology is now evolving at "unprecedented velocity," with models "becoming smarter and more robust every day." App. 205 (Department strategic plan); *see also id.* at 85 (Amodei acknowledging "a smooth, unyielding increase in AI's cognitive capabilities"). They likewise agree that reliable AI technology is critical for the United States to maintain its military superiority and preserve

its national security. *Id.* at 202–07 (Department strategic plan); *id.* at 106 (Amodei on prospect of "unbeatable army" controlled by AI). As explained above, Anthropic definitively rejected usage terms demanded by the United States, in the wake of a significant controversy regarding the Department's use of Claude in an overseas military operation. And two days after Anthropic did so, the United States began offensive military operations in Iran, reportedly using Claude in connection with its strikes. *See* Weisgerber *et al.*, *U.S. Strikes in Middle East Use Anthropic, Hours After Trump Ban*, Wall St. J. (Feb. 28, 2026). Given all this, the Department did not act unconstitutionally by moving to exclude Claude from its supply chain and shift to other AI providers immediately, while providing Anthropic with notice and an opportunity to respond almost immediately thereafter.

Anthropic offers two objections to the Department's claim of urgency. First, it argues that any national-security risk posed by Claude could not have been urgent, since the Department had been deploying Claude—with usage restrictions—for over a year. But as shown above, the Department's use of Claude initially was limited and occurred only through contractors; and Anthropic gradually loosened various usage restrictions as it sought to expand its relationship with the Department. Moreover, the dispute did not come to a head until early 2026, when the Department developed its AI strategic plan and Anthropic definitively rejected the Department's request for an "all lawful uses" authorization. Around the same time, a potentially serious dispute arose regarding Claude's use in connection with one military conflict, which revealed operational concerns as another conflict began to unfold. The timing of these events confirms the Department's legitimate sense of urgency. Second, Anthropic argues that any national-security risk cannot be urgent because the Department permitted use of Claude for six months after the exclusion. But

the Department ordered Claude removed from its supply chain "as soon as practical," with the six-month period established only as an outer bound. App. 80. And it later explained that the outer bound reflected the technical and operational challenges with "remov[ing] the technology from all DoW systems immediately, particularly in the midst of active operations." *Id.* at 216. Again, we are unwilling to second-guess that fact-based judgment of national security. *See Trump v. Hawaii*, 585 U.S. at 704; *HLP*, 561 U.S. at 33.

B

The First Amendment provides that "Congress shall make no law … abridging the freedom of speech." The Supreme Court has held that the Amendment "prohibits government officials from subjecting an individual to retaliatory actions for engaging in protected speech." *Nieves v. Bartlett*, 587 U.S. 391, 398 (2019) (cleaned up). To succeed on such a First Amendment retaliation claim, the plaintiff or petitioner must prove that (1) it engaged in protected speech, (2) the government took materially adverse action against it, and (3) the speech caused the materially adverse action. *See Houston Cmty. Coll. Sys. v. Wilson*, 595 U.S. 468, 477–79 (2022); *Aref v. Lynch*, 833 F.3d 242, 258 (D.C. Cir. 2016).

Anthropic has satisfied the first and second prongs of this test, but not the third. The First Amendment squarely protects Anthropic's advocacy regarding the safe and appropriate use of AI products. Moreover, the Department's exclusion of Claude from its supply chain plainly qualifies as a materially adverse action. However, we can discern no causal connection between the two. Instead, the record makes clear that the Department removed Anthropic from its supply chain not because of its advocacy, but because Anthropic refused to agree to a contract term the Department deemed essential to national security.

Consider the timeline. By its own admission, Anthropic has advocated for use and safety restrictions since its founding: The company describes its own commitment to model safety as "the core of Anthropic's mission," App. 2, and describes its use restrictions as reflecting "the very purpose for which [the] company was founded" and its "deeply held values," *id.* at 10. But throughout 2024 and 2025, the Department never acted adversely to Anthropic. Instead, it included Anthropic in a $200 million AI contract awarded in July 2025, *id.* at 32–33, and it sought to expand its relationship with Anthropic. Additionally, in January 2026, Amodei published a long article calling for "limits" and "safeguards" on the use of AI-powered weapons in democracies. *Id.* at 106, 108. But instead of retaliating against Anthropic because of this advocacy, the Department continued to pursue negotiations. Only when the negotiations broke down did the Department take action: On February 24, the Secretary met with Amodei and demanded that Anthropic accept an "all lawful uses" term by February 27. On February 26, Amodei published a final, public refusal to assent to that term. On February 27, the Secretary announced his intention to exclude Anthropic from the Department's supply chain. And on March 3, he made the formal, written determination required by the Supply Chain Security Act.

Anthropic points to various pungent statements in the Secretary's February 27 social media post. Among other things, the Secretary denounced Anthropic's "sanctimonious rhetoric," "virtue-signaling," and "Silicon Valley ideology." App. 77. Such rhetoric seldom provides a sound basis for judging the lawfulness of federal executive action. *See*, *e.g.*, *Mullin v. Doe*, 146 S. Ct. 2121, 2139 (2026); *Trump v. Hawaii*, 585 U.S. at 700–02. In any event, for all its flourishes, the Secretary's social media post squarely addresses Anthropic's refusal to provide the "all lawful uses" contractual authorization. He described Anthropic's behavior as a

"textbook case of how not *to do business*" with the Pentagon. App. 77 (emphasis added). He reiterated the Department's demand for "full, unrestricted access to Anthropic's models for every lawful purpose in defense of the Republic." *Id.* (cleaned up). And he characterized Anthropic's refusal to provide that access as imposing an unacceptable "veto power over the operational decisions of the United States military." *Id.* The nub of this dispute was contractual, and the First Amendment did not require the Department to continue a contractual relationship that it viewed as creating a national-security risk.

VI

This case raises profoundly difficult questions about the appropriate military uses of an almost unimaginably powerful new technology. The Secretary raises the deeply sobering prospect of overly constrained AI models shutting down unexpectedly and thus causing important military operations to fail. Anthropic raises the deeply sobering prospect of unconstrained AI models hallucinating inappropriate targets for lethal military force. Both possibilities present obvious national-security concerns. But in our Republic, it is the President and the Secretary of War who must determine how best to balance the competing risks. In doing so here, the Secretary did not transgress any limits on his authority under the Supply Chain Security Act or the Constitution. Accordingly, we deny the petitions for review.

So ordered.

KAREN LECRAFT HENDERSON, *Circuit Judge*, dissenting: Whether the Secretary of the Department of War (Secretary) lawfully invoked his statutory powers under the Federal Acquisition Supply Chain Security Act of 2018 (FASCSA) turns on whether Anthropic falls within the statute's definition of a "supply chain risk." 10 U.S.C. § 4713(k)(6). "When Congress takes the trouble to define the terms it uses," *Garland v. Cargill*, 144 S. Ct. 1613, 1627 n.9 (2024) (citation modified), courts should apply them "with rigor," Antonin Scalia & Bryan A. Garner, *Reading Law: The Interpretation of Legal Texts* 227 (2012) (interpretive-direction canon). And here, the Congress has taken great pains to define the type of "supply chain risk" that must exist before the Secretary invokes the sweeping powers FASCSA confers on him. Under the statute, he may exercise his authority to blacklist a procurement source from the Department's supply chains "only after" the Department concludes the source poses a "significant" risk, 10 U.S.C. § 4713(b)(1), that it will:

> sabotage, maliciously introduce unwanted function, extract data, or otherwise manipulate the design, integrity, manufacturing, production, distribution, installation, operation, maintenance, disposition, or retirement of covered articles so as to surveil, deny, disrupt, or otherwise manipulate the function, use, or operation of the covered articles or information stored or transmitted on the covered articles.

10 U.S.C. § 4713(k)(6).

My colleagues do not dispute that whether Anthropic qualifies under this definition depends entirely on the scope of section 4713(k)(6)'s residual clause: "or otherwise manipulate." *See* Maj. Op. 23–26. The Department and Anthropic offer competing definitions of that term. According

to Anthropic, the residual clause uses "manipulate" to denote intentionally subversive acts, carried out through deceptive means. Pet'r Br. 40, n.2; *see, e.g.*, *Manipulate*, Oxford English Dictionary, https://perma.cc/9NF9-ME3B ("To manage, control, or influence in a subtle, devious, or underhand manner."). The Secretary argues, Resp't Br. 42, and the majority agrees, Maj. Op. 24, that the term encompasses much more. To "manipulate" a covered article, in their view, means to "move, arrange, operate, or control [it] by the hands or another body part or by mechanical means," regardless of purpose or motive—for example, how one might "manipulate" a doorknob by turning it or a gas pedal by pressing it down. *Manipulate*, American Heritage Dictionary of the English Language (5th ed. 2022), https://perma.cc/B2FA-CM84 ("She manipulated the lights to get just the effect she wanted."). Both definitions may be linguistically possible but basic canons of construction require us to decide which one fits best within section 4713(k)(6)'s surrounding text, construed as a whole. Because I believe that the context decidedly favors the narrower reading, I respectfully dissent.

Section 4173(k)(6) begins by linking "manipulate" with the opening list of verbs and modified verb phrases: "sabotage, maliciously introduce unwanted function, extract data, or otherwise manipulate" specific aspects of a covered article. The canon of *noscitur a sociis* teaches that "a word is known by the company it keeps." *McDonnell v. United States*, 579 U.S. 550, 569 (2016) (citation modified). If "words [are] grouped in a list," they "should be given related meanings." *Third Nat'l Bank in Nashville v. Impac Ltd., Inc.*, 432 U.S. 312, 322 (1977). Applied here, the correct definition of "manipulate" aligns best with the examples that precede it—all of which connote intentionally subversive and deceptive acts. To "sabotage" means to "ruin, destroy or disable" something "deliberately and maliciously" and "frequently by indirect,"

Sabotage, Oxford English Dictionary (3d ed. 2026), https://perma.cc/E3B4-4JRV, or "underhanded" means, *Sabotage*, American Heritage Dictionary of the English Language, New College Edition (1976). To "maliciously introduce unwanted function [and] extract data" likewise implies an "intention or desire to . . . cause injury." *Malice*, Oxford English Dictionary (3d ed. 2026), https://perma.cc/Z499-HDYV. When we at last reach the "or otherwise manipulate" clause, the narrower meaning—"to manage, control, or influence in a subtle, devious, or underhand manner"—fits well with its antecedents. *Manipulate*, Oxford English Dictionary (3d ed. 2026), https://perma.cc/9NF9-ME3B.

Section 4713(k)(6) uses "manipulate" to round out another verb-object sequence—"so as to surveil, deny, disrupt or otherwise manipulate" an article's function, operation, use or transmission and storage of information. And as with the statute's first string of verbs, the verbs preceding "manipulate" in the second series connote some intentionally hostile or clandestine purpose with respect to a covered article. *Surveil*, Oxford English Dictionary (3d ed. 2026), https://perma.cc/B4UZ-KWWV ("To exercise surveillance over (someone)."); *Surveillance*, Oxford English Dictionary (3d ed. 2026), https://perma.cc/6YR5-H9S7 ("Watch or guard kept over a person . . . often, spying, supervision."); *Disrupt*, American Heritage Dictionary of the English Language (5th ed. 2022), https://perma.cc/7MEA-ZRWQ ("To throw into confusion or disorder" or "[t]o interrupt or impede the progress of"); *Deny*, American Heritage College Dictionary (4th ed. 2007) ("To decline to grant or allow; refuse."). Here again, the more tailored meaning of "manipulate" completes the list of verbs with semantic precision. *See*, *e.g.*, *Manipulate*, Merriam-Webster's Collegiate Dictionary (11th ed. 2003) ("To control

or play upon by artful, unfair, or insidious means esp. to one's own advantage.").

Without expressly invoking the *noscitur a sociis* canon, my colleagues give several semantic reasons to resist its application. Maj. Op. 24. But their arguments follow neither the plain text of section 4713(k)(6) nor well-settled canons. They assert that section 4713(k)(6)'s first parallel verb list is not sufficiently "tight" to qualify for the canon because two of the verbs ("introduce" and "extract") have direct objects in front of them ("unwanted function" and "data"). *Id.* Yet the canon hardly demands such exactitude. *See United States v. Fischer*, 64 F.4th 329, 375 (D.C. Cir. 2023) (Katsas, J., dissenting) (correctly observing that a "listing" is not even a "prerequisite" for the canon) (citation modified), *rev'd*, *Fischer v. United States*, 144 S. Ct. 2176 (2024); *accord* Scalia & Garner, *supra*, at 197 ("An 'association' is all that is required."). A listing that "involves . . . verbs *and* verb phrases" is plainly enough. *Id.* (emphasis added).

My colleagues also assert that the verbs included in section 4713(k)(6)'s lists lack a common quality of harmful or subversive intent. Starting with the statute's first verb sequence, they concede, as they must, that "sabotage" "connotes intentionally hostile acts." Maj. Op. 23. But they conclude that "introduc[ing] unwanted function" carries no such connotation if it is cut off from its modifier "maliciously." *Id.* Respectfully, I disagree with that approach: "[I]t is a fundamental principle of statutory construction (and, indeed, of language itself) that the meaning of a word cannot be determined in isolation, but must be drawn from the context in which it is used." *Reno v. Koray*, 515 U.S. 50, 56 (1995) (citation modified); *Fischer*, 64 F.4th at 365 (Katsas, J., dissenting) ("[W]e do not divorce isolated words and phrases from their statutory context."). They conclude next that the

verb phrase "extract data" connotes no harmful motive at all, Maj. Op. 23, but like its next-door neighbor ("introduce unwanted function"), it is preceded by the adverb "maliciously," 10 U.S.C. § 4713(k)(6). Under the series-qualifier canon, "a prepositive modifier" before a series of verbs or verb phrases "normally applies to the entire series," not only the verb that immediately follows. Scalia & Garner, *supra*, at 147; *id.* at 148 (explaining "*willfully* damage or tamper with" modifies "both damage and tamper with.") (emphasis added). Thus, a library might post a sign saying, "Do not shout, loudly talk on the phone, play music, or otherwise disturb others." The common understanding would be that the rule bans bringing a boom-box into the reading room with the volume turned on high but not listening to music with headphones set at a modest sound level—even though both constitute "playing music." As with the "loudly" modifier, section 4713(k)(6)'s "maliciously" extends beyond its head phrase "introduce unwanted function" to limit "extract data," as well.

With respect to section 4713(k)(6)'s second verb chain, the majority concludes that the verbs are either "ambiguous" as to motive ("surveil" and "disrupt") or "clearly" neutral on that score ("deny"). Maj. Op. 24. But this repeats the same mistake, in my view, of relying on statutory terms in isolation. *Fischer*, 64 F.4th at 363 (Katsas, J., dissenting) ("Perhaps no interpretive fault is more common than the failure . . . to consider the entire text.") (quoting Scalia & Garner, *supra*, at 167). To "surveil" or "deny" something, of course, does not always imply nefarious intent: "The police officer surveilled the suspect's last known whereabouts" or "the cashier denied my credit card for insufficient funds." And "disrupting" something can be unintentional: "The weather disrupted our travel plans." But viewed in their statutory context, the verbs at issue are all directed at deliberately impeding or

eavesdropping on the "function, use, or operation" of a covered article that has entered the federal supply chain. 10 U.S.C. § 4713(k)(6). And once paired with the terms in the first verb-object sequence—*e.g.*, "maliciously introduc[ing] unwanted function" into a covered article to "deny" its use, "sabotag[ing]" the article to "disrupt" its operation, "maliciously . . . extracting data" from the article to "surveil" the information it stores or transmits information—the statute's focus on harmful motive becomes impossible to ignore.

Reading section 4713(k)(6)'s "or otherwise manipulate" clause as limited by the intentionally hostile and deceptive acts described in section 4716(k)(6) becomes even plainer "under the related canon of *ejusdem generis*." *Fischer*, 144 S. Ct. at 2184. That rule declares that, if a statute introduces a list of specifics followed by a catchall phrase, we should construe the general phrase as "controlled and defined" by the examples "that precede it." *Id.* (citation modified). The canon ensures that courts "give effect, if possible, to every clause and word of a statute," *Williams v. Taylor*, 529 U.S. 362, 404 (2000), and reflects "the basic logic that Congress would not go to the trouble of spelling out [a] list" of examples only to tack on a general clause that makes them irrelevant. *Fischer*, 144 S. Ct. at 2185. Faithful application of that principle requires us to read section 4713(k)(6)'s residual clause as targeting a class of interference with a covered article that is not so broad as to read out the statute's numerous examples of deliberately obstructive and surreptitious acts: "sabotag[ing]," "surveil[ing]," "disrupt[ing]," "deny[ing]" and "maliciously" inserting a function into or extracting data from a covered article. And for reasons I have explained, the narrower reading of "manipulate" effortlessly follows that rule.

In adopting the Secretary's neutral definition of "manipulate," today's decision "renders an unnerving amount"

of section 4716(k)(6)'s motive-infused language "mere surplusage." *Fischer*, 144 S. Ct. at 2190. Using the Secretary's sterilized definition, the majority concludes that a contractor immediately poses "not only a 'risk' . . . but a certainty" of "manipulat[ing]" a covered article, Maj. Op. 23, if it is "willing and able to enforce contractual restrictions" on an article's functioning, *id.* at 19, that the Department deems "unreasonably restrictive," J.A. 181. It matters not if the restrictions are ones the Department has already agreed to. And it is immaterial if the contractor enforces the restrictions based on a good faith and legally correct interpretation of them. It does not even seem to matter under the Department's reading if those contractual and technical safeguards prevent the Department from deploying the covered article in a manner that violates federal or constitutional law. The Department here made good on its promise to designate Anthropic a supply chain risk after Anthropic declined the Secretary's ultimatum to replace its use restrictions on Claude with a general provision permitting "all lawful uses." Maj. Op. 9. But suppose the Secretary tells Anthropic's presumed replacement to change its AI-use policies to permit any "functions that the Department deems necessary" or it will share the same fate as Anthropic. *Id.* at 20. According to today's decision, that contractor will have a choice: Agree to the Secretary's demands or risk being designated a national security threat under FASCSA.

I cannot agree that this is the scenario the Congress had in mind when it enacted FASCSA. It enacted the statute in response to calls from the U.S. intelligence community for legislation to meet the threat of "[h]ostile nation state and other bad actors" infiltrating the federal government's information and technology systems through its supply chains. S. Rep. No. 115-408, at 2 (2018). For years, national security agencies had warned that companies "beholden to foreign governments,"

Open Hearing on Worldwide Threats: Hearing Before the S. Select Comm. on Intel., 115th Cong. 64 (2018) (statement of Chris Wray, Dir. of the Fed. Bureau of Investigation), and other malicious actors were introducing compromised products into "[m]any of the technologies the Federal Government relie[d] on for vital, daily functions," S. Rep. No. 115-408, at 2. In their published reports, the agencies described numerous covert security breaches carried out by nefarious actors—and in terms that closely track section 4713(k)(6)'s key terms. *See, e.g.*, Tara Beeny, U.S.-China Econ. & Sec. Rev. Comm'n, *Supply Chain Vulnerabilities from China in U.S. Federal Information and Communications Technology* 23, 27, 34 (Apr. 2018) (describing attacks in which bad actors attempted to "extract proprietary source code," "[]install[ed] unwanted software" onto compromised computers, and "surveil[ed] and manipulate[d] users by hacking . . . embedded [computer] firmware"). Such historical evidence showing how a statute's terms were used pre-enactment sheds light on what sense those words are meant to carry when the Congress writes them into law. *Bostock v. Clayton County*, 140 S. Ct. 1731, 1749–50 (2020). That history supports Anthropic's reading. And it refutes the view that "manipulat[ion]" of a covered article encompasses anything like the conduct that, under today's holding, gives rise to a supply chain risk—that is, a contractor's honest and upfront enforcement of restrictions on a covered article's use disfavored by the government.

For the foregoing reasons, I respectfully dissent.